Compositing Multi-System Infrastructure Pressure at the Census-Subdivision Scale

Data, model and results behind the VISION Canada combined-risk dashboard

Urban Predictive Analytics Lab

UBC Smart Cities · School of Community and Regional Planning, University of British Columbia

Reference implementation: `integrated_dashboard.html` · Figures reproducible via `tools/paper_stats.js` · Reference configuration: SSP2, 2050

ABSTRACT

Infrastructure planning in Canada is organised by sector. Population projection, housing demand, road capacity, water demand and climate exposure are each modelled, funded and governed separately, and each is usually mapped on its own. The practical consequence is that a community can appear unremarkable on five sectoral maps while sitting in the worst decile of all five at once — a condition no single map can express. This report documents the data, model and results behind the combined-risk layer of the VISION Canada platform, which addresses that gap by compositing eight pressure indicators into a single index for every census subdivision (CSD) in Canada over the 2025–2050 horizon under five Shared Socioeconomic Pathways.

The model is deliberately simple: each indicator is percentile-ranked across every subdivision that reports it, and the ranks of the user-selected indicators are averaged under a strict complete-case rule, so a score expresses simultaneous relative position across systems rather than a summation of incommensurable physical units. Seven pre-configured indicator sets ("planning lenses") expose the fact that the choice of what to composite is a policy judgement, not a technical default.

We verify the implementation by independent re-computation, which reproduces the live application's scored counts exactly for all seven lenses. We then report the index's behaviour. Three findings qualify how the output should be read. First, indicator collinearity is severe in places: population growth and water-demand growth correlate at $r = 0.972$, so the water lens largely re-ranks population, whereas population growth and heat intensification are very nearly independent ($r = 0.081$) and therefore carry genuinely additive information. Second, the index is robust in geography but not in magnitude across pathways: the identity of the flagged communities is nearly pathway-invariant, while the population living in them ranges from 17.5 to 32.0 million between SSP3 and SSP5. Third, indicator discretisation interacts with the implementation's positional tie-handling: 64.9% of subdivisions report exactly zero road congestion yet receive ranks spread across 64.9 rank points, shifting individual scores by up to 24.9 points and altering 8–22% of top-quintile membership in the affected lenses. We document each limitation with quantified sensitivity, and set out the interpretive rules the index supports and those it does not.

Keywords: infrastructure interdependency; composite indicators; percentile rank aggregation; census subdivision; Shared Socioeconomic Pathways; CMIP6 downscaling; multi-system risk; planning support systems

1 Introduction

Canadian infrastructure planning is sectorally organised at every level that matters: models are built by sector, capital budgets are allocated by sector, and authority is divided between orders of government by sector. A municipality's water system, arterial network and housing supply are each planned on thirty- to fifty-year asset lives, through different programmes, by different institutions, on different review cycles. Each is also normally visualised on its own map.

This produces a specific and consequential blind spot. Sectoral maps answer the question "where is this system under pressure?" They cannot answer "where is *every* system under pressure at the same time?" The distinction is not academic. A community in the 60th percentile nationally for population growth, housing demand, road congestion and water demand is unexceptional on four maps, yet it faces four capital decisions inside one planning horizon. A community in the 95th percentile on one indicator and the 20th on the rest will dominate a single sectoral map while being, in aggregate, a less complicated planning problem. Sequenced separately, sectoral decisions can produce the familiar failure mode of a subdivision delivered with housing units but inadequate water pressure and an arterial operating at level of service F.

The combined-risk layer documented here is a response to that gap. It composites eight indicators drawn from five independently prepared model families into one index, computed for every census subdivision in Canada, for six five-year steps from 2025 to 2050, under five Shared Socioeconomic Pathways (SSPs). It is implemented as a client-side layer in `integrated_dashboard.html`; all computation described in this report occurs in the browser at request time, against static model outputs.

Our purpose is documentary rather than advocative. Composite indicators are easy to build and easy to over-read: aggregation hides the behaviour of the parts, and equal weighting is a substantive assumption that looks like a neutral default. Accordingly this report gives equal space to the construction of the index (§ 4), the verification of its implementation (§ 5), its empirical behaviour (§ 6), the interpretive claims it can and cannot support (§ 7), and its limitations stated with quantified sensitivity (§ 8). Readers who need only the operational rules for reading a score may begin at § 7.

1.1 Design commitments

Four decisions shape everything that follows, and each is a trade-off rather than an optimum.

1. **Ranks, not units.** The indicators are measured in incommensurable units — percent, degrees Celsius, dwelling counts. Rather than normalising to a common physical scale, which would require defensible exchange rates between a degree of warming and a missing dwelling, each indicator is converted to its percentile position within the set of subdivisions that report it. The index therefore measures *relative standing*, and can say nothing about absolute severity (§ 7.2).
2. **Equal weights.** Selected indicators are averaged without weighting. No empirical basis exists for asserting that, say, congestion matters 1.4 times as much as aging; an unweighted mean states the assumption plainly instead of burying a judgement in coefficients. The cost is that collinear indicators effectively receive double weight (§ 6.3, § 8.1).
3. **Complete cases only.** A subdivision receives a score only if it has data for *every* selected indicator. This preserves the meaning of the composite — genuine simultaneous overlap — at the cost of substantial and unevenly distributed attrition, since one indicator has national coverage of only 23% (§ 4.5).

4. **Explicit lenses.** Because the result depends on which indicators are composited, the interface does not offer a single canonical index. It offers seven named indicator sets, each corresponding to a recognisable planning question, and reports how many subdivisions each set can actually score. Lens choice changes the flagged set materially (§ 6.5).

2 Study area, unit of analysis and computational context

2.1 Unit of analysis

The unit is the census subdivision (CSD), Statistics Canada's standard geography for municipalities and municipal equivalents, on 2021 Census boundaries. The CSD was chosen because it is the finest geography at which all five contributing model families are jointly available, and because it corresponds to the jurisdiction that actually holds local capital authority in most provinces. Subdivisions are identified by their seven-digit CSDUID; the leading two digits give the province or territory, which the application uses for regional aggregation without a separate lookup.

Three different counts of “all subdivisions” appear in this report and must be distinguished, because coverage percentages depend on which denominator is used.

Table 1. Reference counts for the census-subdivision universe. Coverage figures elsewhere in this report use the 5,161 geometry count as the denominator, matching the application's own reporting.

Count	n	Source and meaning
Mapped geometries	5,161	`csd_geojson.json` — every polygon the map draws, including subdivisions with no model output. This is the denominator for the “of 5,161 scored” figure shown on screen.
Modelled subdivisions	4,928	The CSD universe of the UPAL projection datasets (population, housing, water). Excludes subdivisions the projection pipeline could not resolve.
Maximum scorable	5,157	Subdivisions with a heat value — the best-covered indicator, and effectively the whole country. Any lens drawing on the projection datasets is capped near 4,896 instead.

2.2 Temporal frame

The composite is defined at six points: 2025, 2030, 2035, 2040, 2045 and 2050, with 2025 serving as the baseline against which growth indicators and the change view are measured. The horizon is set by the shortest-running input rather than by choice: the population, housing and transport projections terminate at 2050, although the platform's other layers extend to 2100. When the year selector is set beyond 2050 the composite is held at 2050 and the interface states that it has done so, in preference to silently extrapolating or blanking (§ 4.9).

2.3 Computational context

The dashboard is a static document; there is no server-side computation and no database. Model outputs are pre-computed JSON, fetched over HTTP and composited in the browser on each interaction. Two properties of this architecture bear on the methodology. First, every figure on screen is derived from the same in-memory computation, so the map, the ranking table, the scatter plot, the narrative text and the CSV export cannot disagree with one another. Second, because indicator files are cached per URL, changing pathway or lens re-composites from memory rather than re-downloading; the climate files in particular carry all four CMIP6 pathways in one file per province per decade, so switching pathway triggers no new request.

3 Data

Eight indicators are drawn from six files produced by five independent modelling efforts. They differ substantially in provenance, temporal status and spatial coverage, and those differences propagate directly into the composite. Table 2 gives the formal definition and source of each; Table 3 gives measured coverage. The narrative subsections that follow record the provenance caveats that a user of the index needs in order to read it responsibly.

Table 2. The eight indicators: definition, source, unit and temporal status. "Dynamic" indicators respond to the year selector; "fixed" indicators are single snapshots that do not. Sign convention gives the direction in which the raw value indicates greater pressure.

Indicator	Definition	Source file	Unit	Temporal status	Sign
pop Population growth	Percent change in projected population from 2025 to the selected year	nia2_dashboard/population.json UPAL projection	%	Dynamic, 2025–2050	higher
housing Housing demand growth	Percent change in projected total dwelling-unit demand from 2025 to the selected year	nia2_dashboard/housing.json UPAL projection	%	Dynamic, 2025–2050	higher
transport Road congestion	Share of road kilometres modelled at or over capacity in the selected year (level of service E–F)	nia2_dashboard/transport.json UPAL capacity model	%	Dynamic, 2025–2050	higher
undersupply Housing undersupply	Dwelling-unit mismatch, stock minus demand. Negative values denote a shortfall	housing_mismatch/housing_mismatch_data.json	units	**Fixed** at the 2025 snapshot	lower (inverted)
aging Aging population share	Population aged 65 and over as a share of total projected population in the selected year	nia2_dashboard/population.json UPAL projection	%	Dynamic, 2025–2050	higher
heat Extreme-heat intensification	Increase in the annual hottest-day maximum temperature relative to the 1971–2000 baseline, ensemble median	climate_data/<PR>/Hottest_Day_<year>.json CanDCS-M6 (CMIP6)	°C	Dynamic, decadal snapping	higher
water Water-demand growth	Percent change in daily municipal water demand ($\mathrm{ML\ day^{-1}}$) from 2025 to the selected year	water_consumption_2100/csd_data.json Bayesian demand model	%	Dynamic, 2020–2100	higher
infra Road-condition strain	Share of local roads rated poor or very poor by the reporting municipality	nia2_dashboard/ccpi.json Statistics Canada CCPI 2022	%	**Fixed** at 2022	higher

Table 3. Measured indicator coverage at SSP2/2050: the number of subdivisions receiving a percentile rank, as a share of the 5,161 mapped geometries. Coverage is the single most important determinant of how many subdivisions a lens can score. Seven indicators are effectively national; the binding constraint on the platform is road condition, and it is a survey-response limit outside the platform's control.

Indicator	Ranked	Coverage	Binding constraint on coverage
heat	5,157	99.9%	All 13 provinces and territories, all 13 decadal files; no missing values
pop	4,896	94.9%	Projection universe (4,928) less subdivisions with a zero or absent 2025 base
aging	4,895	94.8%	As above; requires a non-zero projected population in the target year
undersupply	4,852	94.0%	Mismatch dataset universe
water	4,788	92.8%	Requires a non-zero 2025 demand baseline
housing	4,607	89.3%	Requires a non-zero 2025 demand base; excludes subdivisions with no modelled demand
transport	4,258	82.5%	Road-network model universe; subdivisions with no modelled network are absent
infra	1,200	23.3%	**Survey response:** 1,298 of ~5,160 subdivisions reported the relevant CCPI cells; 1,200 have a usable road-condition value

3.1 Population and aging

Population is the platform's anchor dataset, projecting all 4,928 subdivisions across five pathways in five-year steps to 2050, with an age–sex breakdown from which the 65-and-over share is derived. National 2050 growth relative to 2025 differs sharply across pathways — SSP1 +23%, SSP2 +20%, SSP3 0%, SSP4 +14%, SSP5 +40% — and these divergent totals are the mechanism behind the magnitude sensitivity reported in § 6.4. Population enters the composite as an indicator, but is also loaded unconditionally regardless of lens, because the narrative and KPI panels report the headcount living in the flagged subdivisions.

Two properties of this dataset warrant note. Projected populations are continuous rather than integer, so small subdivisions carry fractional headcounts; and because growth enters as a *percentage* change, subdivisions with very small 2025 populations can post extreme percentages from small absolute changes. We test whether this distorts the index in § 6.6 and find that it does not dominate in aggregate, though it does shape the extreme tail of the ranking.

3.2 Housing demand and undersupply

Two distinct housing indicators enter the composite and should not be conflated. **Housing demand growth** is forward-looking and pathway-dependent: the change in modelled unit demand between 2025 and the selected year, covering all thirteen provinces and territories. **Housing undersupply** is a different quantity from a different file: a stock-versus-demand balance in dwelling units, in which negative values denote a shortfall. Across the 4,852 subdivisions reporting it, the balance ranges from a shortfall of 412,152 units to a surplus of 92,348, with a median of +23 units — that is, the median Canadian subdivision is in slight nominal surplus while the aggregate national shortfall is concentrated in a small number of large urban subdivisions.

Critically, the undersupply indicator is a **fixed 2025 snapshot**. It does not respond to the year selector. A user comparing 2030 and 2050 under a lens containing undersupply is varying only the other indicators; the undersupply contribution is constant. The interface states this beside the indicator checkbox, and it is the reason the housing-crunch lens shows a nearly identical top quintile under mid-rank as under positional tie-handling (§ 8.3) — its inputs are continuous and static.

3.3 Road congestion and road condition

Congestion is expressed as the share of a subdivision's modelled road kilometres operating at or over capacity. The source documentation is explicit that no trip-generation, travel-demand or origin–destination data exists anywhere in the contributing pipeline: peak-hour volume is a *proxy* for travel demand, and inter-municipal commuting flows are not represented. Congestion in a dormitory suburb generated by commuting into an adjacent city is therefore not attributed across the boundary. The indicator should be read as "modelled local network loading" rather than as an observed traffic measurement.

Road condition is drawn from the Statistics Canada Canadian Core Public Infrastructure Survey (CCPI) 2022 and is the share of local roads a municipality rated poor or very poor. Its limitations are inherited wholesale from the survey. Coverage is 1,298 subdivisions (~25%) because that is who answered those cells; it is a fixed 2022 snapshot with no scenario or time dimension; and it covers the "local roads" asset class only. Coverage here is genuine survey non-response, not processing loss, and it cannot be remedied downstream.

3.4 Extreme-heat intensification

The heat indicator is the projected increase in the annual hottest-day maximum temperature relative to a 1971–2000 baseline, taken from the CanDCS-M6 statistically downscaled CMIP6 ensemble. Each file holds, per subdivision, the 10th, 50th and 90th ensemble percentiles for both the absolute value and the change from baseline, under each of four CMIP6 pathways; the composite uses the **ensemble median change** (`d<scenario>_p50`), discarding the ensemble spread. Files exist at thirteen decadal points from 1951 to 2071.

Coverage is complete: all 13 provinces and territories are present, at all 13 decadal points, with identical schema and no missing values — 5,157 subdivisions in total, which is effectively the entire country and makes heat the best-covered indicator in the set. Two other constraints are material, however. Only the ensemble median is used, so the spread across climate models is discarded, and only one heat index is represented: humidity, heat-wave duration and night-time minima are all absent, although each is material to heat mortality. And values are stored at 0.1 °C resolution, yielding only 23 distinct values across 5,157 subdivisions at SSP2/2050; this heavy discretisation interacts with tie-handling as documented in § 8.3.

3.5 Water demand

Water enters as the percent change in daily municipal water demand between 2025 and the selected year, from a national Bayesian demand model covering all five pathways from 2020 to 2100. Each record stores a triple of served population, demand in megalitres per day, and per-capita demand in litres per day; the composite reads the second element.

TWO CAUTIONS ON THE WATER INDICATOR

It is demand, not scarcity. The indicator measures growth in withdrawal requirement. It says nothing about whether supply exists to meet it. A subdivision with rapidly growing demand and abundant renewable supply ranks identically to one with the same growth and no headroom. Supply-side stress is modelled elsewhere on the platform and is deliberately not composited here.

It is largely a function of population. Because demand is modelled as served population multiplied by a per-capita rate, its percentage change tracks population change closely — measured at $r = 0.972$ (§ 6.3). A lens containing both population and water therefore weights population growth approximately twice.

4 Model

The composite is computed in five steps: extraction of a raw value per indicator per subdivision, sign normalisation, conversion to percentile rank, application of the complete-case filter over the selected indicator set, and an unweighted mean. Figure 1 shows the pipeline; the subsections give each step formally.

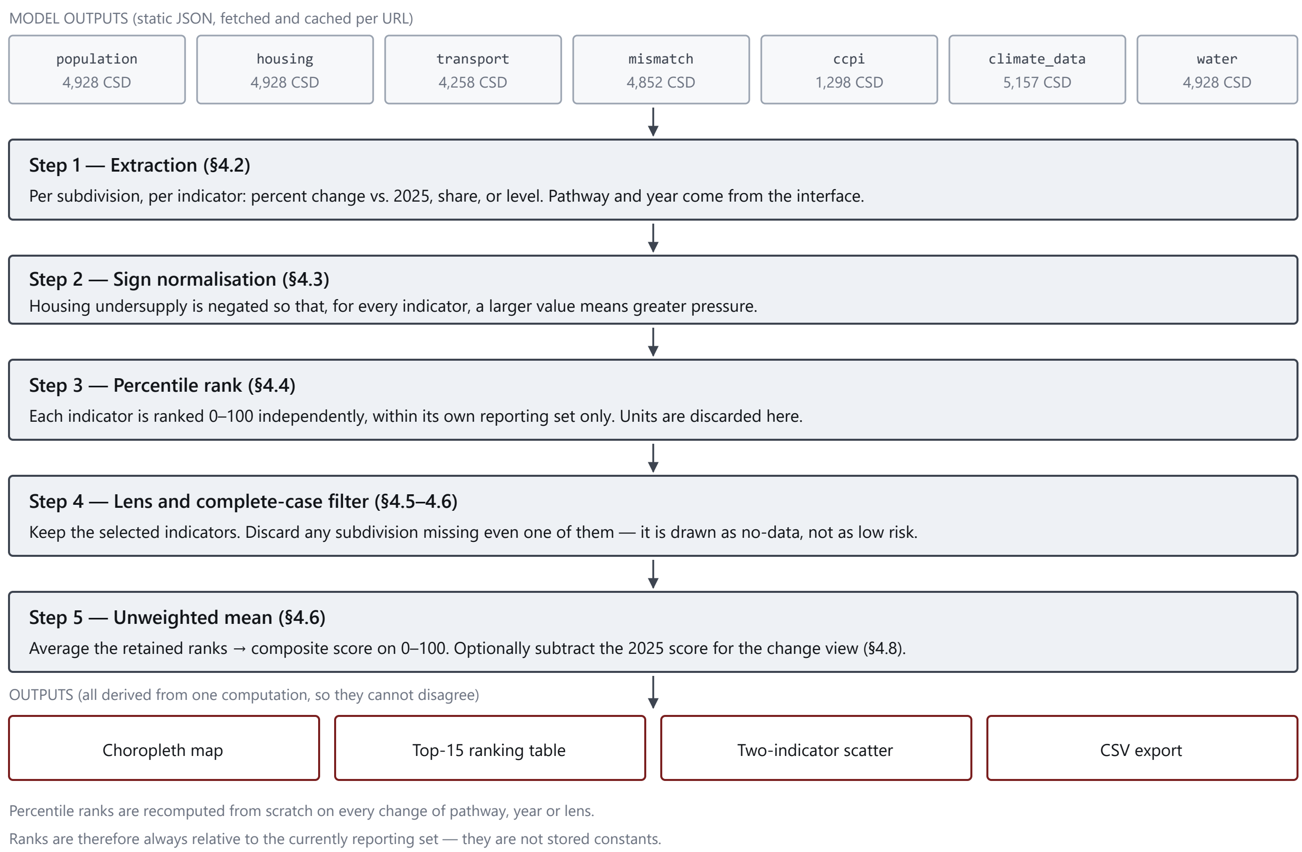

Figure 1. The combined-risk pipeline as implemented in `integrated_dashboard.html`. All five steps execute client-side on each interaction. The critical structural feature is that ranking (Step 3) precedes filtering (Step 4): an indicator's percentile ranks are computed across everyone who reports *that indicator*, before the complete-case rule removes subdivisions missing any *selected* indicator. The consequences are discussed in § 4.7.

4.1 Notation

Let C be the set of census subdivisions and K = {pop, housing, transport, undersupply, aging, heat, water, infra} the indicator set. For pathway s and year t, let $x_{i,k}(s,t)$ denote the raw value of indicator k for subdivision i, which may be undefined. Let $L \subseteq K$ be the active lens, the indicator subset the user has selected.

4.2 Step 1: indicator extraction

Five indicators are growth rates against the 2025 baseline, computed identically:

$$x_{i,k}(s,t) = 100 \cdot (\ v_{i,k}(s,t) - v_{i,k}(s,\ 2025)\ )\ /\ v_{i,k}(s,\ 2025) \tag{1}$$

for $k \in$ {pop, housing, water}, where v is the projected level. Undefined when the 2025 base is zero or absent.

The remaining indicators are taken as levels or shares rather than changes:

- **transport** — the modelled share of road kilometres at or over capacity in year t, read directly.
- **aging** — the 65-and-over share, computed as $pop_{65+}(s,t)\ /\ pop(s,t) \times 100$.
- **heat** — the ensemble-median change in hottest-day temperature relative to the 1971–2000 baseline, at the nearest available decadal file (§ 4.9).
- **undersupply** and **infra** — read from their fixed snapshots, independent of t.

Note the asymmetry this creates, which is intentional but easily missed: population, housing and water contribute *trajectories*, while congestion, aging, heat, undersupply and condition contribute *states*. A subdivision can therefore rank high because it is changing fast, or because it is already in a poor condition, and the composite does not distinguish the two.

4.3 Step 2: sign normalisation

Ranking requires that a larger value mean greater pressure for every indicator. Seven indicators already satisfy this. Housing undersupply does not: it is signed, with negative values denoting shortfall. It is therefore negated before ranking:

$$\hat{x}_{i,k} = -x_{i,k} \ \textit{if}\ k = \textit{undersupply}, \qquad x_{i,k}\ \textit{otherwise} \tag{2}$$

This is the only transformation applied to raw values. No logarithm, winsorisation or outlier treatment is applied, because the subsequent rank transformation is invariant to any monotone transformation and therefore makes them redundant.

4.4 Step 3: percentile rank

For each indicator independently, let $R_k = \{i \in C : \hat{x}_{i,k} \text{ is defined}\}$ be its reporting set, and let $n_k = |R_k|$. Sort R_k ascending by $\hat{x}$ and let $\pi_k(i) \in \{0, \ldots, n_k - 1\}$ be the resulting zero-based position. Then

$$r_{i,k} = 100 \cdot \pi_k(i)\ /\ (\ n_k - 1\ ) \tag{3}$$

so that the lowest-pressure subdivision receives 0 and the highest receives 100.

Three properties of Equation 3 must be stated explicitly, because each shapes the results.

It is a positional rank, not a proportion below. The extremes are always exactly 0 and 100 regardless of the underlying distribution, and the ranks are uniformly distributed by construction. The transformation therefore discards all information about magnitude and about the shape of the distribution: the gap between the 1st- and 2nd-ranked subdivision may be a hundred times the gap between the 50th and 51st, and the composite cannot see the difference.

Ties are not averaged. Subdivisions with identical raw values receive *different* ranks, determined by their position in the sort — which, for a stable sort over object-key order, is effectively arbitrary. Where an indicator is heavily discretised this is consequential: 2,762 of 4,258 subdivisions (64.9%) report exactly zero congestion at 2050, and those tied subdivisions are spread across 64.9 rank points. § 8.3 quantifies the effect on results and states the remedy.

Ranks are relative to the indicator's own reporting set, not to the scored set. Road condition is ranked among only the 1,200 subdivisions that answered the CCPI question. A road-condition rank of 90 therefore means "in the worst decile *of the municipalities that responded*", not of Canada. Because ranking precedes filtering, this remains true even when the lens reduces the scored set further (§ 4.7). Six of the eight indicators are close enough to national coverage that the distinction is immaterial for them; for road condition it is not.

4.5 Step 4: the complete-case rule

A subdivision is scored only if it has a rank for every indicator in the active lens:

$$S_L = \{\ i \in C : r_{i,k} \text{ is defined for all } k \in L\ \} \tag{4}$$

Subdivisions outside S_L receive no score, are drawn in the dedicated no-data grey, and are counted in the legend and narrative as unscored. This is a substantive choice with a clear rationale and a real cost.

The rationale is that the index claims to measure *simultaneous* pressure. Were missing indicators simply skipped — averaging whatever is available — a subdivision reporting only one indicator, on which it happened to rank 99th percentile, would score 99 and outrank a subdivision genuinely in the 85th percentile on all eight. The composite would then rank subdivisions partly by how little is known about them, which inverts the intended meaning.

The cost is attrition that is institutionally structured rather than random. Because coverage is limited chiefly by which municipalities answered a 2022 survey, the surviving set is not a random sample of Canada. In the full eight-indicator lens only 1,165 of 5,161 subdivisions (22.6%) are scorable, and they are necessarily a subset of the CCPI respondents — which is why that lens's leading communities cluster in interior British Columbia (§ 6.2). That result is a statement about data availability at least as much as about risk.

4.6 Step 5: aggregation and the planning lenses

The composite is the unweighted mean of the retained ranks:

$$Score_i(L, s, t) = (\ 1\ /\ |L|\ ) \cdot \textstyle\sum_{k \in L} r_{i,k}(s, t)\ , \qquad i \in S_L \tag{5}$$

Seven lenses are pre-configured. They are not alternative estimators of one underlying quantity; they are different questions, and they return materially different answers (§ 6.5).

Table 4. The seven planning lenses, with the question each is intended to answer and the coverage each achieves at SSP2/2050. Coverage is a property of the lens, not of the pathway or year, and is the first thing a user should read.

Lens	Indicators	Planning question	Scored	Coverage
Growth pressure	pop + housing + transport	Where is everything rising together?	4,087	79.2%
Housing crunch	pop + undersupply	Which growing places are already short of stock?	4,820	93.4%
Aging & mobility	aging + transport	Which older communities face congestion?	4,168	80.8%
Heat & growth	pop + heat	Where does growth meet intensifying extreme heat?	4,892	94.8%
Infrastructure strain	transport + infra	Where is congestion compounded by poor road condition?	1,176	22.8%
Water demand pressure	pop + water	Where does growth drive water demand hardest?	4,788	92.8%
Full multi-system	all eight	Where does every measured system converge?	1,165	22.6%

Any other combination of the eight indicators may be selected directly; the lenses are entry points, not a closed set. Coverage for an arbitrary combination is bounded above by its weakest constituent indicator in Table 3.

4.7 Why ranking precedes filtering

The order of Steps 3 and 4 in Figure 1 has an effect worth isolating, because it is the most common source of misreading. Ranks are computed over each indicator's full reporting set, and only afterwards is the scored set narrowed to subdivisions complete on the lens. Consequently a subdivision's population rank is its standing among ~4,896 subdivisions, even when the active lens can only score 1,165 of them. The score is thus a composite of *national* standings, evaluated on a *restricted* sample.

The alternative — re-ranking within the scored set after filtering — would make each score internally consistent with its sample, and would force every lens's scores to span the full 0–100 range. We note the trade-off rather than resolving it: the present choice keeps a given subdivision's population rank stable as the user toggles other indicators, which makes the interface legible, at the cost that the mean of national ranks over a restricted set does not span 0–100. This is directly visible in the results: the full lens's top-quintile threshold is 70.2 and its maximum is 88.6, not 100 (§ 6.2), because no subdivision in that restricted sample is extreme nationally on all eight indicators at once.

4.8 The change view

A second view colours the change in composite score since the baseline year:

$$\Delta Score_i(L, s, t) = Score_i(L, s, t) - Score_i(L, s, 2025) \tag{6}$$

The baseline score is a full recomputation at 2025 using the same pathway, lens and rank procedure, so ΔScore is a change in *relative national standing*, not a change in physical conditions. A positive value means a subdivision has moved up the national ranking; it may do so while its absolute conditions improve, if others improve faster. The change view is symmetric about zero on a diverging colour scale so that "no change" is always the neutral midpoint, and it is disabled when the selected year is itself 2025. Note that the two fixed-snapshot indicators (undersupply, infra) contribute identical ranks at both ends and therefore damp ΔScore toward zero in any lens containing them.

4.9 Pathway and temporal alignment

The composite must reconcile three input families with different scenario vocabularies and different time grids. Three alignment rules do this, and each is surfaced in the interface rather than applied silently.

Pathway pairing. The socioeconomic pathway selected by the user drives the population, housing, transport and water models directly. The climate archive, however, is indexed by CMIP6 forcing scenario, so the heat indicator requires a pairing:

Table 5. SSP-to-CMIP6 pairing used for the heat indicator, with the measured heat response at 2050. The ordered progression of median warming across pathways is a useful internal coherence check.

Socioeconomic pathway	CMIP6 scenario	File key	CSDs	Median ΔT	Range
SSP1 — Sustainability	SSP1-2.6	`d126_p50`	5,157	+2.60 °C	+1.5 to +3.5
SSP2 — Middle of the road	SSP2-4.5	`d245_p50`	5,157	+3.40 °C	+2.1 to +4.4
SSP3 — Regional rivalry	SSP3-7.0	`d370_p50`	5,157	+3.90 °C	+2.5 to +4.8
SSP4 — Inequality	**none**	—	0	—	—
SSP5 — Fossil-fuelled development	SSP5-8.5	`d585_p50`	5,157	+4.80 °C	+3.0 to +6.4

SSP4 has no Tier-1 CMIP6 counterpart in the downscaled archive. Rather than substituting a neighbouring pathway, the heat indicator returns no data under SSP4, so any lens containing heat scores zero subdivisions on that pathway. This is deliberate: an unavailable pairing is reported as unavailable.

Climate-year snapping. Climate files exist at decadal points, while the composite runs on five-year steps. The heat indicator uses the nearest available file:

$$t_{climate} = argmin_{y \in Y} \mid y - t \mid , \qquad Y = \{1951, 1961, \ldots, 2061, 2071\} \tag{7}$$

So a composite at 2050 reads the 2051 file, at 2045 the 2041 file, and at 2025 the 2021 file. The maximum temporal offset is five years, small relative to the decadal signal being measured, but it means the heat indicator moves in steps as the year selector sweeps, while the other indicators move smoothly.

Horizon clamping. The year selector spans 2020–2100 for the platform's other layers, but the composite's inputs stop at 2050. Outside 2025–2050 the composite is held at the nearer endpoint, and a badge and rail note state which year is actually being reported, so that the map, ranking, scatter and export are never silently mislabelled.

5 Verification

The composite is computed in browser JavaScript, embedded in a presentation layer, with no unit tests. To establish that the algorithm described in § 4 is the algorithm actually executing, we re-implemented it independently (`tools/paper_stats.js`) as a Node script reading the same files over the filesystem rather than over HTTP, and compared its scored counts against the counts the live application reports in its own narrative panel. Agreement is exact for all seven lenses.

Table 6. Verification at SSP2/2050. The live-application counts were read from the dashboard's insight panel under headless browser automation; the independent counts come from the offline re-implementation. Exact agreement across all seven lenses indicates that the documented algorithm matches the executing one.

Lens	Live application	Independent re-implementation	Agreement
Growth pressure	4,087	4,087	✓
Housing crunch	4,820	4,820	✓
Aging & mobility	4,168	4,168	✓
Heat & growth	4,892	4,892	✓
Infrastructure strain	1,176	1,176	✓
Water demand pressure	4,788	4,788	✓
Full multi-system	1,165	1,165	✓

This verifies internal consistency between documentation and implementation. It is not a validation of the underlying projections, which is outside the scope of this report, and it does not test the correctness of the input models.

6 Results

All results are reported at the platform's reference configuration — SSP2, 2050 — except where a pathway comparison is the subject. §§ 6.1–6.2 describe what the index produces; §§ 6.3–6.6 test how it behaves.

6.1 Per-lens results

Table 7. Composite results by lens at SSP2/2050. The top quintile is defined as subdivisions at or above the 80th percentile of the scored distribution, so its *count* is fixed at roughly a fifth by construction; what varies between lenses is *which* subdivisions qualify and how many people live in them.

Lens	k	Scored	80th pct. threshold	Top quintile	Population in top quintile	Pop. share	Leading provinces (count)
Growth pressure	3	4,087	74.2	818	24,282,560	50.8%	QC 164, AB 163, SK 151, BC 147
Housing crunch	2	4,820	73.5	964	26,835,885	56.0%	QC 237, BC 185, ON 141, SK 127
Aging & mobility	2	4,168	64.7	834	14,448,865	30.2%	QC 198, BC 196, ON 107, SK 95
Heat & growth	2	4,892	70.1	979	14,354,147	29.9%	SK 292, BC 163, QC 162, AB 156
Infrastructure strain	2	1,176	76.5	236	17,537,212	50.7%	ON 71, BC 54, SK 32, AB 32
Water demand pressure	2	4,788	80.0	958	17,536,187	36.6%	SK 245, QC 195, AB 179, BC 146
Full multi-system	8	1,165	70.2	233	11,689,952	33.8%	AB 69, ON 57, BC 49, SK 35

Population figures are projected 2050 headcounts under SSP2 for the subdivisions in question. The infrastructure-strain and full lenses are confined to the CCPI respondent set, so their provincial distributions reflect survey response as much as risk; the other five lenses are effectively national.

Two patterns in Table 7 are worth drawing out. First, the population share held by the top quintile varies nearly twofold across lenses, from 29.9% (heat & growth) to 56.0% (housing crunch), even though the quintile count is structurally fixed. The housing-crunch and infrastructure lenses concentrate on large urban subdivisions — roughly half the scored population sits in a fifth of the subdivisions — while the aging and heat lenses distribute across smaller communities. The composite is therefore informative about *who* is exposed, not merely *where*.

Second, the aging & mobility lens has a markedly lower threshold (64.7) than the others. This follows from indicator opposition rather than from lower risk: aging share and congestion are *negatively* correlated ($r = -0.397$, § 6.3), so averaging them compresses scores toward the middle and few subdivisions rank high on both. A lens built from opposed indicators will always produce a narrower score distribution, and its top quintile identifies the genuinely unusual combination rather than the extreme of either indicator.

6.2 Leading communities and the coverage artefact

Table 8. Five highest-scoring subdivisions under three lenses at SSP2/2050. The composition of these lists is itself a result: the full lens is restricted to the CCPI respondent set, and its ranking is correspondingly regional, whereas the two national lenses above it range freely across the country.

Lens	Highest-scoring subdivisions (score /100)
Growth pressure	Picture Butte, AB (99.5) · Skowkale, BC (99.1) · Kwawkwawapilt 6, BC (98.8) · Skwah 4, BC (98.8) · Calmar, AB (98.7)
Heat & growth	Gainsborough, SK (98.5) · Cupar, SK (98.1) · Alida, SK (97.6) · Maryfield, SK (97.4) · Carstairs, AB (97.2)
Full multi-system	Chase, BC (88.6) · Peachland, BC (88.3) · Creston, BC (86.3) · Armstrong, BC (86.2) · Summerland, BC (86.1)

The full-lens result illustrates § 4.5 and § 4.7 together. Its five leading subdivisions are all in interior British Columbia, and its maximum score is 88.6 rather than approaching 100. Both facts are consequences of construction rather than discoveries about Canadian geography: the scorable set is the 1,165 subdivisions that answered the CCPI road-condition cells, and within that restricted sample no subdivision is nationally extreme on all eight indicators simultaneously. A user who reads this list as "Canada's most infrastructure-stressed communities" will be wrong; the correct reading is "among the 23% of subdivisions where all eight systems can be measured, these rank highest on the average of their national standings."

The heat & growth lens makes the same point from the opposite direction. Now that the climate indicator covers all 13 provinces and territories, that lens scores 4,892 subdivisions and its ranking is led by small Saskatchewan communities — Gainsborough, Cupar, Alida, Maryfield — where high projected warming coincides with growth. Saskatchewan holds 292 of the 979 top-quintile subdivisions, the largest provincial share of any lens in Table 7. That geography was invisible while the indicator was restricted to Alberta, British Columbia and Ontario, which is a useful reminder that the apparent findings of a composite index are bounded by the extent of its inputs before they are bounded by anything substantive.

Conversely, the presence of several very small subdivisions at the top of the growth-pressure lens — including three First Nations reserves — reflects percentage growth computed on small 2025 bases (§ 3.1). § 6.6 tests whether this dominates the index and finds that it does not, but it does shape the extreme tail, and rankings truncated at the top five will over-represent small communities.

6.3 Indicator collinearity

Because the composite is an unweighted mean, its behaviour depends on how independent its constituents are. Two indicators measuring the same latent quantity contribute twice to the average while appearing to be one factor each. We therefore computed Pearson correlations between percentile ranks over all pairwise-complete subdivisions.

Table 9. Correlation matrix of indicator percentile ranks at SSP2/2050, computed pairwise-complete. Shading runs from strong positive (dark) through neutral (light grey) to negative (blue). Cell counts vary because coverage differs; the smallest pairings are those involving road condition (n ≈ 1,180–1,200), while every other pair exceeds n = 4,000.

	pop	housing	transp.	unders.	aging	heat	water	infra
pop	—	0.910	0.326	0.267	−0.098	0.081	0.972	0.204
housing	0.910	—	0.327	0.340	−0.277	0.085	0.896	0.206
transport	0.326	0.327	—	0.236	−0.397	0.215	0.332	0.189
undersupply	0.267	0.340	0.236	—	−0.291	0.046	0.301	0.171
aging	−0.098	−0.277	−0.397	−0.291	—	−0.269	−0.128	−0.066
heat	0.081	0.085	0.215	0.046	−0.269	—	0.099	0.008
water	0.972	0.896	0.332	0.301	−0.128	0.099	—	0.201
infra	0.204	0.206	0.189	0.171	−0.066	0.008	0.201	—

The matrix separates the indicator set into three groups with quite different implications for the composite.

A near-redundant growth cluster. Population, housing demand and water demand form a tight bundle: pop×water r = 0.972, pop×housing 0.910, housing×water 0.896. This is mechanical rather than empirical — housing demand and water demand are both modelled substantially as functions of projected population — and it has a direct consequence for interpretation. The water-demand-pressure lens (pop + water) is, at r = 0.972, close to ranking population growth against itself; its apparent two-system agreement is largely an artefact of shared inputs. The growth-pressure lens is better behaved because congestion is only weakly related to the other two ($r \approx 0.33$), but two of its three indicators still measure much the same thing.

Near-independent pairings. Heat is very weakly related to the growth cluster: pop×heat r = 0.081, housing×heat 0.085, water×heat 0.099. Where a community is growing and where warming is most severe are close to unrelated questions in Canada, so the heat & growth lens combines two effectively independent dimensions and carries the most additional information per indicator of any lens in Table 4 — it identifies communities that could not be found by examining either indicator alone. Road condition is likewise weakly related to everything ($|r| \leq 0.21$, and only 0.008 against heat), consistent with asset condition being a legacy of past investment rather than a function of projected growth or climate.

Opposed pairings. Aging is negatively related to congestion (−0.397), undersupply (−0.291), housing demand (−0.277) and heat (−0.269): older communities are, on average, the slower-growing, less congested and less rapidly-warming ones. Averaging opposed indicators is legitimate — the aging & mobility lens is asking precisely where the unusual conjunction occurs — but it compresses the score distribution (§ 6.1) and means that a mid-range score in that lens is the norm rather than a sign of moderate risk. The one modest positive relationship involving heat is with congestion (0.215), which most likely reflects that both are elevated in southern, more urbanised latitudes rather than any causal link.

6.4 Pathway sensitivity

Because the composite ranks rather than measures, changing the socioeconomic pathway scales all subdivisions' underlying trajectories together and leaves their relative order largely intact. The consequence is a sharp separation between the geography of the result, which is nearly pathway-invariant, and its magnitude, which is not.

Table 10. Pathway sensitivity of the growth-pressure lens at 2050. The scored count, the quintile count and the leading subdivision are effectively constant, while the population living in the top quintile spans 14.4 million between the lowest and highest pathway.

Pathway	National 2050 growth	Scored	Threshold	Top quintile	Population in top quintile	Leading subdivision
SSP1 — Sustainability	+23%	4,087	74.0	818	26,552,502	Picture Butte, AB (99.5)
SSP2 — Middle of the road	+20%	4,087	74.2	818	24,282,560	Picture Butte, AB (99.5)
SSP3 — Regional rivalry	0%	4,087	72.2	818	17,549,645	Picture Butte, AB (99.3)
SSP4 — Inequality	+14%	4,089	73.4	818	22,303,636	Picture Butte, AB (99.4)
SSP5 — Fossil-fuelled	+40%	4,088	74.4	818	31,975,855	Picture Butte, AB (99.5)

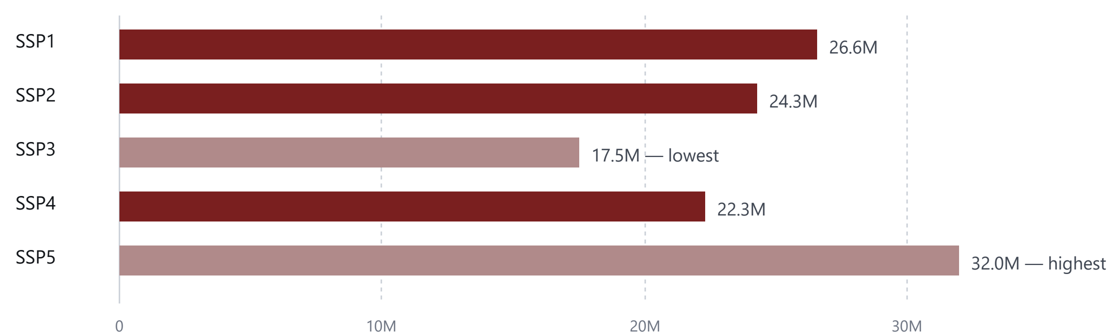

Projected 2050 population of the 818 top-quintile subdivisions, growth-pressure lens. The subdivisions are almost the same set in every pathway.

Figure 2. Pathway sensitivity is a question of magnitude, not of geography. The same ~818 subdivisions are flagged under every pathway, but the population exposed to the flagged conditions ranges from 17.5 million under SSP3 to 32.0 million under SSP5 — a factor of 1.8. Pathway choice should therefore be understood as setting the stakes of a planning problem whose location is already robust.

This separation is the single most useful robustness property of the index and the one most easily misread in the opposite direction. Because Picture Butte leads the ranking under all five pathways and the threshold moves by only 2.2 points, a user might conclude that pathway choice is immaterial. It is immaterial to *where to look* and highly material to *how much capacity to build*: the difference between SSP3 and SSP5 is 14.4 million people in the same communities. Conversely, a user should not read pathway invariance as validation — it is a property of rank aggregation, and would appear even if the underlying projections were substantially wrong in level, provided they were wrong proportionately.

6.5 Lens sensitivity

If the lenses were interchangeable views of one latent risk, they would flag substantially the same communities. They do not. Table 11 reports the Jaccard index of top-quintile membership between each pair.

Table 11. Agreement between lenses on top-quintile membership at SSP2/2050, as the Jaccard index (shared ÷ union). Values are low throughout: the highest pair agrees on barely half its combined membership, and most pairs share less than a fifth.

Lens pair	Shared	Jaccard	Interpretation
Growth pressure × Water demand pressure	613	0.527	Highest agreement, driven by the pop×water collinearity of §6.3
Heat & growth × Water demand pressure	607	0.456	Both share the population indicator, on nearly the same footprint
Housing crunch × Water demand pressure	537	0.388	Shared population indicator
Growth pressure × Housing crunch	493	0.382	Shared population indicator
Growth pressure × Heat & growth	444	0.328	Near-independent second indicator moves a third of the set
Infrastructure strain × Full multi-system	113	0.317	Shared restriction to the CCPI respondent set
Growth pressure × Aging & mobility	304	0.226	Opposed indicators (§6.3) pull the sets apart
Growth pressure × Full multi-system	208	0.247	Adding five indicators retains only a quarter of the combined set
Heat & growth × Infrastructure strain	80	0.070	Lowest agreement of any pair

Selected pairs shown; the complete 21-pair matrix is produced by `tools/paper_stats.js`, section 7. Jaccard values range from 0.070 to 0.527 across all pairs, with a median of 0.187.

No pair exceeds 0.53, and the median pair shares under a fifth of its combined membership. The practical implication is that the lens selector is not a display preference but the most consequential choice a user makes — more consequential than the pathway (§ 6.4), which barely changes membership at all. Presenting a single default composite would therefore have concealed the most important source of variation in the output, which is why the interface requires an explicit lens and reports its coverage alongside its result.

6.6 Community size composition

Percentage-change indicators computed on small denominators can produce extreme values from trivial absolute changes, raising the concern that the index systematically elevates very small communities (§ 3.1, § 6.2). We tested this by comparing the size-class composition of each lens's top quintile against its scored set.

Table 12. Size-class composition of the scored set versus the top quintile, at SSP2/2050. Size classes follow the platform's own thresholds, derived from 2021 Census population (Very Small <1,000; Small 1,000–4,999; Medium 5,000–24,999; Large 25,000–99,999; Very Large 100,000+).

	Growth pressure (4,087 scored)		Full multi-system (1,165 scored)	
Size class	**Scored share**	**Top-quintile share**	**Scored share**	**Top-quintile share**
Very Small	48.8%	28.7%	18.8%	9.4%
Small	32.7%	28.4%	40.9%	33.9%
Medium	13.7%	28.7%	29.2%	37.8%
Large	3.0%	9.7%	7.2%	13.3%
Very Large	1.4%	4.0%	3.9%	5.6%
Median 2025 population	**1,003**	**3,509**	**3,414**	**7,177**

The concern is not borne out in aggregate; the bias runs the other way. In the growth-pressure lens, Very Small subdivisions are 48.8% of the scored set but only 28.7% of the top quintile, while Medium subdivisions more than double their share (13.7% to 28.7%) and Large subdivisions triple theirs (3.0% to 9.7%). The median top-quintile subdivision holds 3,509 people against 1,003 for the median scored subdivision — 3.5 times larger. The same direction holds in the full lens.

The reason is the complete-case rule interacting with congestion coverage: very small subdivisions frequently have no modelled road network and are therefore excluded before scoring, and those that survive tend to rank low on congestion. The small-denominator effect is real but confined to the extreme tail — it explains the specific identities in the top five of Table 8 without distorting the quintile as a whole. Users reading a truncated top-*n* list should nonetheless be aware of it, and the platform's CSV export exists partly so that the full distribution, rather than the visible top 15, can be examined.

6.7 Worked example

Table 13 traces the full computation for two subdivisions at SSP2/2050, chosen to illustrate both the arithmetic and the complete-case rule.

Table 13. Worked computation for Kelowna and Chilliwack, British Columbia, at SSP2/2050. Both are fast-growing southern-interior cities with nearly identical growth profiles; they diverge only in whether the municipality answered the CCPI road-condition question.

Indicator	Kelowna raw	rank	Chilliwack raw	rank	Note
pop	+46.6%	98.6	+42.2%	97.8	Both in the top 3% nationally
housing	+48.8%	97.4	+48.3%	97.3	Tracks population closely
transport	6.1%	81.0	44.6%	93.8	Chilliwack far more congested
undersupply	−12,591	99.3	−7,323	98.7	Severe shortfall in both
aging	18.9%	41.0	17.2%	31.5	Below-median 65+ share
heat	+3.7 °C	90.5	+3.3 °C	45.0	0.4 °C spans 45 rank points
water	+37.6%	98.7	+33.5%	98.0	Mirrors population, per §6.3
infra	3%	39.9	**not reported**	—	Chilliwack did not answer this CCPI cell
Growth pressure `(98.6+97.4+81.0)/3`		**92.4**		**96.3**	Both scored; Chilliwack higher on congestion
Full multi-system `mean of all eight`		**80.8**		**unscored**	Chilliwack is drawn grey: one missing indicator disqualifies it

Three lessons follow from this single pair. **Averaging dilutes.** Kelowna ranks above the 97th percentile on four indicators, yet its full-lens score is 80.8, because two indicators on which it is unremarkable (aging 41.0, road condition 39.9) pull the mean down. A high composite becomes progressively harder to achieve as indicators are added, which is why full-lens scores are lower than narrow-lens scores throughout Table 7 and why scores are not comparable across lenses (§ 7.2). **Grey is not safe.** Chilliwack is unscored in the full lens purely because its municipality did not answer one survey question; on the seven indicators it does report, it is at least as pressured as Kelowna. **Rank hides magnitude.** Kelowna's heat rank of 90.5 and Chilliwack's of 45.0 correspond to +3.7 °C and +3.3 °C — a 0.4 °C difference that spans 45 rank points, because the underlying distribution is narrow and heavily discretised into 23 distinct values (§ 3.4). The reverse holds too: two communities separated by 45 rank points on this indicator differ by less than half a degree.

7 Interpretation guidance

7.1 What a score means

A composite score is the average national percentile standing of a subdivision across the selected indicators, conditional on it reporting all of them. A score of 80 in a three-indicator lens means the subdivision sits, on average, in the worst fifth of Canada on three systems *at once* — not that any single system has crossed a threshold, and not that conditions are 80% as bad as they could be.

The distinction that gives the index its purpose is between a high score and a high single-indicator value. A subdivision at the 99th percentile on congestion and the 30th on everything else scores 53 in a three-indicator lens and will not appear in the top quintile; one at the 82nd percentile on all three scores 82 and will. The index is deliberately built to surface the second case, which sectoral maps cannot show, at the cost of de-emphasising the first, which they show well. The two views are complementary, and the platform retains its single-sector layers for that reason.

> **READING ORDER FOR A DEFENSIBLE INTERPRETATION**
>
> **1. Read the coverage before the colour.** A lens scoring 22.6% of subdivisions supports very different claims than one scoring 94.8%. Coverage is reported next to every result.
> **2. Establish that grey means unknown.** Unscored subdivisions may be under more pressure than any scored one; they are missing an input, not passing a test.
> **3. Treat the ranking as a shortlist, not a priority order.** Score differences of a few points are within the sensitivity documented in § 8.3 and should not be used to order candidates.
> **4. Return to the raw values before acting.** The map popup and the CSV export give the raw value and the rank for every contributing indicator, which is where the substantive planning content lies.
> **5. Vary the lens deliberately.** Because lens agreement is low (§ 6.5), a community appearing in several lenses is a materially stronger finding than one appearing in the lens that happened to be open.

7.2 What a score does not mean

- **Scores are not comparable between lenses.** A score of 85 in a two-indicator lens and 85 in the eight-indicator lens describe different things; adding indicators mechanically compresses scores toward the mean (§ 6.7). Only compare subdivisions within one lens, pathway and year.
- **Scores are not comparable to physical thresholds.** The index has no units and no engineering meaning. It cannot indicate whether a treatment plant will fail, when a road reaches capacity, or how many dwellings are needed. Those questions belong to the sectoral models, which the composite is designed to point toward, not to replace.
- **The absence of a score is not a low score.** This bears repeating because the visual encoding is easy to misread: the no-data grey and the low-value end of the colour ramp are distinct by design, but only the legend distinguishes them.
- **Differences of a few points are not meaningful.** Given the tie-handling sensitivity of § 8.3 — mean absolute score shifts of 3.4 to 5.2 points in the affected lenses — only substantial gaps should be treated as real.
- **A high score is not an allocation decision.** The index measures pressure, not need, cost-effectiveness, equity, adaptive capacity or fiscal capacity. A wealthy subdivision with reserves and a fiscally constrained one may score identically. Nothing in the model represents ability to respond.
- **The composite says nothing about causation.** Correlated pressures may share a driver, or may be coincident. The index locates coincidence; it does not explain it.

8 Limitations

The limitations below are ordered by how much they should change a reader's confidence. The first three are properties of this implementation and are quantified; the remainder are inherited from the inputs or intrinsic to composite indicators generally.

8.1 Collinearity produces implicit weighting

Equal weighting is stated as a design commitment (§ 1.1), but it delivers equal weighting only if the indicators are independent. They are not (§ 6.3). Because water demand is modelled largely as a function of population, the water-demand-pressure lens assigns population an effective weight approaching 2 while appearing to balance two systems; the growth-pressure lens gives the population–housing dimension roughly two-thirds weight against one-third for congestion. The nominal weights in Equation 5 therefore do not describe the effective ones. We report the correlation matrix in full rather than adjusting the weights, because any correction — principal components, inverse-correlation weighting, or dropping indicators — would substitute an opaque statistical judgement for a transparent arithmetic one. Users should read Table 9 as part of the specification of any lens they rely on.

8.2 Coverage-driven selection

The complete-case rule (§ 4.5) makes the scored set a function of data availability, and availability is structured rather than random: it depends principally on which municipalities answered the relevant cells of a 2022 survey. Seven of the eight indicators now reach effectively national coverage, so this limitation has narrowed to a single indicator — but that indicator still governs the two lenses that include it. The full-indicator lens analyses 22.6% of subdivisions and its results are regionally determined before any pressure is measured (§ 6.2). The eight-indicator result should therefore be read as a demonstration of the method on the subset where all inputs exist, not as a national assessment. The six lenses that exclude road condition carry the defensible national claims; this inverts the intuition that more indicators give a better answer.

8.3 Positional tie-handling on discretised indicators

Equation 3 assigns ranks by sort position without averaging ties (§ 4.4). Three indicators are heavily discretised, so tied blocks are large:

Table 14. Tie structure of the discretised indicators at SSP2/2050. The final column gives the range of percentile ranks spanned by the single largest block of identical raw values — ranks that are assigned by sort position rather than by any property of the subdivision.

Indicator	n	Distinct values	Largest tied block	Rank points spanned	Tied value
transport	4,258	1,029	2,762 (64.9%)	64.9	0% of km over capacity
infra	1,200	88	434 (36.2%)	36.2	0% of roads rated poor
heat	5,157	23	729 (14.1%)	14.1	+3.5 °C

Two subdivisions reporting identical zero congestion can therefore receive congestion ranks differing by up to 64.9 points. To quantify the effect on results we recomputed every lens with mid-rank tie-handling, in which tied values share the mean of their positions — the standard treatment — and compared outcomes.

Table 15. Sensitivity to tie-handling. Jaccard agreement of top-quintile membership, and score displacement, between the implemented positional ranking and standard mid-rank ranking, at SSP2/2050.

Lens	Top quintile	Membership retained	Jaccard	Max |Δscore|	Mean |Δscore|
Aging & mobility	834	729	0.776	16.2	5.09
Growth pressure	818	726	0.798	11.7	3.37
Full multi-system	233	217	0.871	6.8	1.48
Infrastructure strain	236	223	0.896	**24.9**	5.18
Heat & growth	979	937	0.918	3.5	1.14
Housing crunch	964	961	0.994	1.2	0.06
Water demand pressure	958	958	1.000	2.8	0.12

The effect is confined but not negligible. Lenses built only from continuous indicators are unaffected: water demand pressure is identical (Jaccard 1.000) and housing crunch nearly so. Lenses containing a discretised indicator lose 8–22% of top-quintile membership, with individual scores displaced by up to 24.9 points in the infrastructure lens and a mean displacement of 1.1–5.2 points. Neither ranking is more "correct" in principle — both are defensible conventions — but mid-ranking is the standard treatment and does not assign different standings to subdivisions with identical measurements. We record this as the clearest candidate for revision (§ 9), and in the interim it sets the resolution below which score differences should not be interpreted (§ 7.2).

8.4 Static indicators in a temporal composite

Two of eight indicators do not respond to the year selector: housing undersupply is fixed at 2025 and road condition at 2022. In a lens containing them, part of the composite is constant while the rest evolves, so a comparison between 2030 and 2050 varies only some of the indicators. This also damps the change view (§ 4.8), since static indicators contribute zero to ΔScore by construction. The interface labels both indicators as fixed, but the composite arithmetic gives no signal that a score mixes projected and historical quantities.

8.5 Proxy and scope limitations inherited from inputs

- **Congestion is modelled, not observed, and is local.** No trip-generation or origin–destination data exists in the pipeline; peak-hour volume proxies travel demand, and inter-municipal commuting is unrepresented (§ 3.3). Congestion generated by commuting into a neighbouring city is not attributed across the boundary.
- **Water is demand, not scarcity.** Growth in withdrawal requirement is measured without reference to available supply (§ 3.5).
- **Heat is national but is one index and one ensemble member.** Coverage is complete across all 13 provinces and territories, but only the hottest-day ensemble median is used — ensemble spread, humidity, heat-wave duration and night-time minima are all discarded, though each bears on heat mortality (§ 3.4).
- **Road condition is one asset class, self-reported, in one year.** Local roads only, as rated by the municipality, in 2022 (§ 3.3).

- **No adaptive or fiscal capacity is represented.** The index measures pressure only; nothing in it reflects reserves, borrowing capacity, institutional capability or existing programmes.

8.6 Intrinsic limitations

- **Rank aggregation discards magnitude.** By construction, the top-ranked subdivision scores 100 whether it exceeds the second by a trivial or an enormous margin (§ 4.4). The index cannot distinguish a nation of uniform mild pressure from one with a few extreme outliers.
- **The quintile is definitional, not diagnostic.** Roughly 20% of scored subdivisions are in the "top quintile" in every configuration (§ 6.1). The count is fixed by the statistic; only membership is informative. There is no configuration in which the index reports that Canada is doing well.
- **Modifiable areal unit and ecological inference.** Results are conditional on CSD boundaries, which vary enormously in area and population; a single value per subdivision conceals all intra-municipal variation, so a score cannot be attributed to any neighbourhood or household within it.
- **Uncertainty is not propagated.** Point estimates enter and a point estimate leaves. The climate ensemble's 10th and 90th percentiles are available in the source files but unused, and the projection families carry no published intervals in the formats consumed here. The index therefore has no confidence bounds.

9 Candidate revisions

The following changes would address specific limitations above. They are listed in order of benefit relative to effort; none has been applied, so the figures in this report describe the implementation as it currently stands.

1. **Adopt mid-rank tie-handling** in Equation 3. This removes an arbitrary source of variation worth up to 24.9 score points (§ 8.3) and is a localised change to one function. Published scores would shift, so it should be versioned explicitly.
2. **Report effective weights alongside each lens.** Deriving each lens's implied weighting from the correlation structure of its constituents, and displaying it, would make the collinearity of § 8.1 visible at the point of use rather than only in this report.
3. **Address road condition as the single remaining coverage constraint.** It is now the only indicator far from national coverage, and it alone halves what the full lens can see (§ 8.2). Options are to source a condition measure with wider coverage, to model condition from asset age and climate where the survey is silent, or to treat the indicator as an optional overlay reported separately from the composite rather than as an eighth averaged term.
4. **Carry the climate ensemble spread through the composite.** The 10th and 90th percentiles already exist in the source files; propagating them would allow the heat contribution, at minimum, to be reported with an interval.
5. **Distinguish trajectory from state indicators in the interface.** The composite currently averages rates of change and levels without marking the difference (§ 4.2); labelling them, or offering separate sub-scores, would clarify whether a high score reflects rapid change or an already-poor condition.

6. **Add a supply-side or capacity indicator.** Every current indicator measures demand or condition. A capacity or headroom measure would let the index distinguish pressure from pressure that cannot be absorbed.

10 Reproducibility

Every figure in this report can be regenerated from the repository. The implementation under description is `integrated_dashboard.html`; the verification and analysis script is `tools/paper_stats.js`, which is an independent re-implementation rather than a wrapper around the dashboard code, so agreement between them is meaningful.

Table 16. Where each result in this report comes from.

Result	Source	Notes
Tables 7, 8 — per-lens results	`paper_stats.js` §1	SSP2/2050
Table 3 — coverage	`paper_stats.js` §2	Denominator 5,161 geometries
Table 9 — correlations	`paper_stats.js` §3	Pearson on ranks, pairwise complete
Table 10, Figure 2 — pathway sensitivity	`paper_stats.js` §4	Recomputes all five pathways
Table 5 — heat by pathway	`paper_stats.js` §5	Climate file 2051
Table 12 — size composition	`paper_stats.js` §6	Platform size classes
Table 11 — lens agreement	`paper_stats.js` §7	All 21 pairs computed
Tables 14, 15 — tie handling	`paper_stats.js` §8	Positional vs. mid-rank
Table 6 — verification	Dashboard vs. script	Live counts read under browser automation
Table 13 — worked example	Dashboard map popup	Reproducible by clicking either subdivision

To regenerate the analysis, with the data tree populated as described in the repository README:

```
node --max-old-space-size=4096 tools/paper_stats.js
```

To rebuild this document as a PDF:

```
node tools/build_paper.js
```

The composite reads six files: `nia2_dashboard/population.json`, `nia2_dashboard/housing.json`, `nia2_dashboard/transport.json`, `nia2_dashboard/ccpi.json`, `housing_mismatch/housing_mismatch_data.json`, `water_consumption_2100/csd_data.json`, plus `climate_data/<PR>/Hottest_Day_<year>.json` for all 13 provinces and territories, and the boundary file `csd_geojson.json`. Model outputs are produced by the Urban Predictive Analytics Lab from Statistics Canada 2021 Census geography and projections, and from CanDCS-M6 downscaled CMIP6 climate data. Statistics Canada inputs are used under the Open Government Licence — Canada.

11 Conclusion

The combined-risk layer answers a question that sectoral infrastructure analysis structurally cannot: where do independent pressures land on the same community inside the same capital horizon? Its method is intentionally transparent — rank, filter, average — because a composite whose behaviour cannot be traced by hand is not usable as planning evidence.

The empirical results establish three things about how the output should be used. The index is robust in the geography it identifies and sensitive in the magnitude it implies: the same communities are flagged under every socioeconomic pathway, while the population exposed varies by 14.4 million between the lowest and highest (§ 6.4). Its constituents are unevenly independent, so some lenses aggregate genuinely new information — population and heat are effectively orthogonal — while others largely re-rank a single latent quantity, population growth (§ 6.3). And the choice of lens, not the choice of pathway, is the dominant source of variation in which communities are flagged, with no two lenses agreeing on more than 53% of their combined top quintile (§ 6.5).

The most consequential limitation is not statistical but institutional. The index's coverage is set by which municipalities answered a 2022 survey, and the complete-case rule converts that gap into silence: the eight-indicator lens can speak for only 22.6% of Canadian subdivisions, and the grey on the map marks communities about which the platform can say nothing rather than communities that are safe. Seven of the eight indicators now reach effectively the whole country, which leaves road condition as the single binding constraint. Closing it would do more for the reliability of multi-system infrastructure assessment in Canada than any refinement of the aggregation method.

References and sources

1. Statistics Canada. *Census Profile, 2021 Census of Population* and 2021 Census geographic attribute files. Census subdivision boundaries and population base. Used under the Open Government Licence — Canada.
2. Statistics Canada. *Canadian Core Public Infrastructure Survey (CCPI), 2022*. Non-anonymised project extract. Source of the local-road condition indicator; 1,298 reporting census subdivisions.
3. Canadian Centre for Climate Services. *CanDCS-M6: Canadian Downscaled Climate Scenarios — Multivariate, CMIP6*. Statistically downscaled CMIP6 ensemble; source of the hottest-day temperature indicator, 1971–2000 reference period.
4. O'Neill, B. C., Kriegler, E., Ebi, K. L., et al. (2017). The roads ahead: Narratives for shared socioeconomic pathways describing world futures in the 21st century. *Global Environmental Change*, 42, 169–180.
5. Riahi, K., van Vuuren, D. P., Kriegler, E., et al. (2017). The Shared Socioeconomic Pathways and their energy, land use, and greenhouse gas emissions implications. *Global Environmental Change*, 42, 153–168.
6. O'Neill, B. C., Tebaldi, C., van Vuuren, D. P., et al. (2016). The Scenario Model Intercomparison Project (ScenarioMIP) for CMIP6. *Geoscientific Model Development*, 9(9), 3461–3482. Basis for the Tier-1 scenario set underlying the pairing in Table 5.
7. OECD & European Commission Joint Research Centre (2008). *Handbook on Constructing Composite Indicators: Methodology and User Guide*. OECD Publishing. Reference for normalisation, weighting and robustness practice in composite indicators.

8. Transportation Research Board (2010). *Highway Capacity Manual*. National Academies of Sciences, Engineering, and Medicine. Basis for the level-of-service and capacity conventions used by the contributing road model.
9. Urban Predictive Analytics Lab. *VISION Canada 2.0 platform documentation* (repository `README.md`) and the population, housing, transport and water demand model outputs consumed by this layer.

Urban Predictive Analytics Lab · UBC Smart Cities · School of Community and Regional Planning, University of British Columbia. This report documents `integrated_dashboard.html` as implemented at the time of writing; the figures it quotes are reproducible via `tools/paper_stats.js`.