Extended Abstract

Title: Beyond the Density Paradox: Dual Thresholds and Pareto Optimization of the Urban Climate Dilemma (Korea, 2014-2023)

Keywords: density paradox; compact city; mitigation-adaptation trade-off; mobility networks; transport CO_2; panel threshold regression; Pareto frontier (knee point); COVID-19 mobility shock

1. Motivation: resilient urban growth under recurrent shocks

Urban growth in the 2020s is increasingly shaped by compound shocks. Cities face repeated floods and heatwaves while simultaneously being asked to decarbonise—especially through the transport sector. Compact-city policies are widely promoted because higher density, mixed land use, connected street networks, and transit accessibility can reduce car dependence and per-capita transport emissions. Yet the same intensification can magnify hazard impacts by increasing imperviousness (runoff and flood depth), intensifying heat storage and reducing ventilation, and concentrating people and assets along mobility corridors that structure daily exposure.

This creates a decision problem rather than a simple "co-benefit" narrative: compactness can generate mitigation gains and adaptation losses at the same time. Two conceptual gaps limit evidence-based planning. First, most empirical studies treat the compactness-outcome relationship as linear (or weakly quadratic), leaving unanswered the central planning question: at what point do marginal mitigation benefits diminish while hazard impacts accelerate? Second, the literature often conflates physical limits with social optima. Even if further densification remains physically feasible, it may already be socially inefficient once marginal benefits flatten and marginal hazard damages steepen. These gaps motivate a dual-threshold framework that separates:
(i) a structural threshold (D*struct): an econometrically identified break in marginal effects; and
(ii) a Pareto knee (D*pareto): a Pareto-efficient stopping point on the mitigation-adaptation frontier.

2. Research questions and contributions

We address five linked questions:
RQ1 (measurement/validity): Can a multidimensional compactness index be constructed in a time-comparable way and grounded in planning-relevant physical proxies (e.g., floor-area ratio and imperviousness)?

RQ2 (structural non-linearity): Do mitigation benefits exhibit diminishing returns while hazard impacts accelerate beyond a compactness regime—implying a structural threshold?
RQ3 (social efficiency): Does the Pareto knee coincide with the structural threshold? If not, what does the gap imply for planning and resilient growth management?
RQ4 (heterogeneity): Do thresholds and gaps differ systematically across metropolitan and non-metropolitan systems?
RQ5 (mechanism/identification): Are hazard impacts driven by short-run activity flows or by built-form "stock" mechanisms embedded in the urban fabric?

The paper makes four contributions. First, it builds a time-consistent Korean Compactness Index (K-CI) aligned with the "3Ds + accessibility" paradigm and explicitly incorporates network connectivity and transit accessibility—key channels linking urban form, mobility systems, and risk exposure. Second, it estimates non-linear compactness effects using municipality- and year-fixed-effects panel threshold regression, reporting confidence regions and allowing for multiple thresholds, with inference robust to cross-sectional dependence. Third, it translates econometric response functions into a decision-support tool by recovering Pareto frontiers and detecting the knee point; this distinguishes physical tipping regimes from socially efficient stopping points. Fourth, it leverages the 2020–2022 COVID period as a quasi-experimental mobility shock to assess whether compactness gradients persist when activity flows collapse, strengthening mechanism interpretation in line with Cities-at-Risk concerns about mobility networks and urban exposure.

3. Data, variables, and measurement design

We compile a balanced municipality-by-year panel for South Korea (si/gun/gu; 2014–2023; N≈229, T=10). To avoid dilution by non-urban land, compactness measures are computed using urbanised-area denominators. Monetary damage variables are deflated to constant prices.

3.1 Compactness exposure: the Korean Compactness Index (K-CI)

K-CI aggregates four dimensions:

• Density: population density, employment density, and average floor-area ratio (development intensity);

• Diversity: land-use mix measured by an entropy index;

• Connectivity: intersection density as a proxy for network permeability and route choice;

• Accessibility: transit-oriented accessibility (station-area land share), with bus-stop density as a robustness proxy.

To ensure intertemporal comparability, we estimate the index using pooled PCA across all years so that factor loadings are fixed. Because index construction can be a natural target, we conduct systematic sensitivity checks—PCA weights versus equal weights and dimension-balanced weights—verifying that empirical conclusions are not driven by a single weighting scheme. For policy translation, we ground-truth K-CI against planning-relevant physical proxies such as floor-area ratio and imperviousness and report a transparent conversion table that maps changes in K-CI into interpretable physical magnitudes.

3.2 Outcomes: mitigation and adaptation impacts

Mitigation outcome (Y1). Our primary mitigation measure is per-capita road-transport CO_2 emissions (tCO_2 per person), cross-validated using per-capita vehicle-kilometres travelled (VKT) and, where available, satellite-based emissions estimates. Consistent findings across measures are interpreted as evidence that the estimated relationships reflect mobility-system mechanisms rather than inventory artefacts.

Adaptation outcomes (Y2). We focus on realised impacts rather than abstract "risk scores":

• Flood impacts: real per-capita flood damages (currency per person), with flood-event counts used to test sensitivity to heavy tails and reporting;

• Heat impacts: heat-related illness incidence per 100,000 residents.

A standard concern is that observed damages reflect hazard severity rather than vulnerability. We therefore condition explicitly on hazard triggers—extreme-rainfall days (≥80 mm/day) and heatwave days (≥33°C)—and include baseline geographic exposure (e.g., elevation and lowland share). Additional controls capture adaptive capacity, socioeconomic vulnerability, and asset exposure (e.g., GRDP per capita, ageing structure, land values, and fiscal variables). This design aims to estimate the net effect of compactness on impacts under comparable hazard conditions, rather than conflating compactness with climate variability.

4. Empirical strategy I: structural threshold estimation with robust inference

We estimate non-linear effects using a municipality- and year-fixed-effects panel threshold model in which compactness impacts differ across regimes separated by an unknown threshold γ. We test for threshold existence and, critically, allow for the possibility of multiple thresholds. Threshold confidence regions are obtained via bootstrap procedures.

Because municipalities are embedded in commuting systems and regional climate regimes, cross-sectional dependence is likely. We diagnose dependence (e.g., Pesaran-type tests) and base inference on Driscoll-Kraay standard errors robust to general forms of cross-sectional dependence and serial correlation. This addresses a common reviewer concern that threshold findings may be fragile in spatially connected panels.

We also adopt outcome-specific robustness treatments that reviewers routinely request:

• Flood damages often exhibit heavy tails and many zeros. We therefore report results using (i) log(1 + damage) specifications, (ii) winsorisation, and (iii) event-count outcomes as complementary measures.

• For heat, we test sensitivity to alternative hazard definitions (e.g., days above alternative thresholds and annual maximum temperature) to ensure results are not an artefact of a particular cutoff.

The core structural estimand is D*struct: the compactness level at which (a) marginal mitigation benefits flatten and/or (b) marginal hazard impacts steepen. We interpret D*struct as a physical tipping regime rather than a normative optimum.

5. Empirical strategy II: quasi-experimental mobility-shock mechanism check (COVID-19)

Endogeneity is the central identification challenge: compactness may evolve with economic growth, asset values, and policy choices, and hazard exposure may shape development patterns. Our baseline fixed-effects design absorbs time-invariant locational features (e.g., geography) and common time shocks (e.g., national policies), while hazard triggers and exposure controls address climate variability and baseline risk.

To further interrogate mechanisms in a way that speaks directly to “Cities at Risk” and mobility networks, we exploit the 2020–2022 COVID period as an exogenous shock to activity flows. If compactness-related hazard impacts were driven primarily by short-run mobility intensity (flow), we should observe attenuation when activity collapses. Persistence of compactness gradients during COVID is consistent with built-form “stock” mechanisms such as imperviousness and heat storage embedded in the urban fabric. We implement interaction specifications between compactness regimes and COVID-period indicators and interpret regime-specific contrasts as mechanism evidence rather than a standalone causal claim.

Recognising potential COVID confounders (e.g., reporting changes for illness), we complement this test by controlling for healthcare-capacity proxies where available and by focusing on structural patterns (threshold locations and regime differences) that are less sensitive to level shifts.

6. From econometrics to decision support: Pareto frontier, the knee point, and weight sensitivity

Structural thresholds identify breakpoints in marginal effects, but planning is ultimately a multi-objective optimisation problem: reducing transport emissions while limiting hazard impacts. We therefore translate the estimated response functions into a two-objective decision space.

6.1 Pareto frontiers

Holding controls at representative values (and, in sensitivity analysis, at different vulnerability profiles), we simulate predicted outcomes across the observed compactness range to generate a smooth mapping from compactness to (mitigation, adaptation) outcomes. We recover:

• hazard-specific Pareto frontiers for floods and heat; and

• a composite adaptation-impact index (standardised flood and heat impacts aggregated with transparent weights) to summarise multi-hazard trade-offs.

6.2 Knee detection and robustness to weights

We identify D*pareto using curvature-based knee detection on each frontier, interpreting the knee as the point where additional mitigation gains become small relative to the acceleration of adaptation costs. Because composite indices can embed normative choices, we explicitly report weight sensitivity: we vary weights across plausible ranges (including equal weights and hazard-prioritising weights) and test whether the knee location is stable. Our main inference does not rely on a single weighting choice: hazard-specific frontiers provide a weight-free benchmark, and stable knee locations across weights strengthen the claim that the "efficient compactness range" is a structural feature of the system.

We then compare D*pareto to D*struct. The difference (D*struct − D*pareto) is interpreted as a policy-relevant gap between physical limits and social optima: a range where densification may remain physically feasible but is already socially inefficient once hazard damages are internalised.

7. Heterogeneity, inequality, and operational zoning logic

We test heterogeneity in thresholds and gaps across metropolitan versus non-metropolitan systems and across baseline geographic exposure and social vulnerability. These analyses speak to environmental justice concerns: if the efficient compactness range is narrower for vulnerable municipalities, uniform densification policies may inadvertently amplify inequality in climate impacts.

Finally, we translate the dual thresholds into an operational zoning logic:
• Safe zone ($D < D*pareto$): densification is likely to yield net climate benefits; emphasis on transit-oriented development and network efficiency.
• Warning zone ($D*pareto \leq D < D*struct$): densification can be feasible but should be conditional on bundled adaptation investments (stormwater retention, permeable surfaces, open-space protection, urban greening, and heat mitigation), including resilience upgrades for mobility networks (e.g., protecting critical links and improving redundancy).
• Danger zone ($D \geq D*struct$): additional compaction is likely counterproductive without major structural adaptation; down-zoning or strict resilience conditionality is warranted.

8. Expected implications

Preliminary diagnostics indicate meaningful non-linearities consistent with diminishing marginal mitigation returns and accelerating hazard impacts beyond certain compactness regimes. The dual-threshold framework is expected to reveal that the Pareto knee may lie below the structural threshold, implying that socially efficient stopping points can precede physical tipping regimes. By separating "how far we can densify" from "how far we should densify", the paper offers a decision-ready framework for resilient urban growth under recurrent climate shocks and highlights how thresholds vary with mobility-system structure, exposure, and vulnerability.