Antifragile Urban Mobility System: Dynamic Target-Tracking Model with Adaptive Parameters

September 25, 2025

1 Abstract

This report presents a dynamic convergence model for antifragile urban mobility systems, derived from grounded theory analysis of 123 studies (2019–2024). The model quantifies how disruptions trigger adaptive gains, achieving an Antifragility Factor (AF) exceeding context-adaptive thresholds (statistically determined city-specific thresholds). It combines a core dynamical equation with operationalized indicators, calibration, sensitivity analysis, and back-testing, ensuring EU compliance for Horizon Europe deliverables. Key contributions: traceable KPIs, a reproducibility pack, and a transport equilibrium annex. Threshold is NOT a standard and it is a project choice based on statistical significance above noise/seasonal variation, policy-relevant improvement that justifies intervention costs and alignment with EU resilience benchmarks for infrastructure. The threshold can be calibrated per city based on local variability.

This convergence model provides a foundation for future development of a full multi-layer equilibrium framework.

2 Indicator Table & Data Dictionary

All indicators are normalised on a rolling 30-day window:

$$\hat{x} = \frac{x - x_{\min}}{x_{\max} - x_{\min}}$$

Unless noted, variables within $\bar{E}(x)$ are their normalised forms; hats are omitted for readability.

Variable	Scale	Indicator Definition	Data Source	Unit	Aggregation Window
M(t)	Individual / Network	Mobility throughput: avg. multi-modal trips per capita per hour	Smart-card logs; CDR; GPS probes	trips h^{-1} cap^{-1}	Daily average
R(t)	Network	Redundancy ratio: dormant links / active links (pre-event)	OpenStreetMap; traffic sensors	–	Weekly snapshot
A(t)	System	Adaptation velocity: ∂(modal-share diversity)/∂t	Household surveys	$\% \cdot \text{day}^{-1}$	Bi-weekly
C(t)	System	Network entropy of OD matrix (normalized to current topology)	CDR; Bluetooth sensors	nats (normalized)	Daily
D_{eff}(t)	Event	Disturbance (domain/recency/source-weighted)	D2.1 event catalogue + operator/official logs	–	Hourly → daily aggregate
S(t)	System	Stress: Σ(vol/cap)·speed_drop	Loop detectors	–	Hourly
Q(t)	Equity	Stratified Theil–T of accessibility minutes	Census; GTFS	–	Monthly
P(t)	Demand-management	Composite index in [0,1] from time-varying fares via min–max by mode/time, aggregated with ridership weights.	Operator APIs, open fares	index (0–1)	Daily index
B(t)	Behavioral	Compliance and adaptation rate: weighted avg of route guidance compliance, mode shift willingness, and temporal flexibility	App usage data, surveys	index (0–1)	Weekly
E_energy(t)	Energy	EV-grid resilience score	Grid SCADA	–	Daily
E_ICT(t)	ICT	Telecom network uptime during event	Operator logs	%	Event-specific
I(t)	Cross-layer	Synergy index: normalized sum of inverse transfer distances between transport modes	Multilayer graph	–	Weekly

Clarifications and Measurement Notes

- **Disturbance vs. Stress.** Define $D_{\text{eff}}(t)$ as the *external shock magnitude* (event characteristic, input) and $S(t)$ as the *internal congestion response* (network outcome, output). The same event can generate both, but they capture different constructs; practical overlap should be minimal.
- **Inequality metric $Q(t)$.** We consider Theil's T (useful for within/between-group decomposition) and the Gini coefficient (overall inequality). Both range from 0 (equality) to 1 (inequality); either is acceptable if used consistently.
- **Energy/ICT observability and fallbacks.** EV-grid status (SCADA), and ICT uptime often require operator access. When unavailable, use validated proxies or set the indicator to a conservative default (e.g., $E_{\text{energy}} = 1.0$ for "normal") and flag the assumption.
- **Multi-timescale updates.** Fast-moving indices $(P(t), S(t))$ update daily; slow-moving structure ($Q(t)$, infrastructure) updates annually. Multiple time-scales are handled via indicator-specific κ values in the convergence update.

- **Entropy under topology constraints.** With fixed topology, $C(t)$ measures how flows distribute over available routes. During a disruption: *initially* $C \downarrow$ (flows concentrate), *adaptation* $C \uparrow$ (alternatives discovered), *antifragile outcome* C exceeds pre-event levels (more distributed flows). Topology bounds feasible entropy; realised flows determine $C(t)$.

- $P(t)$ **(participation) measurement.** Ridership weights are refreshed weekly/monthly; the daily $P(t)$ uses the most recent weights. Participation is measured via AFC (smart cards/mobile apps) covering $> 90\%$ of trips; quarterly manual surveys address residual gaps.

- **SCADA scope and metrics (if available).** SCADA = Supervisory Control and Data Acquisition. Monitored systems: EV charging network, depot charging, tram/metro supply, traffic-signal power. Metrics: voltage stability ($\pm 5\%$), power availability (% uptime), peak-load/capacity (utilisation), outage duration/frequency, and a power quality index (PQI). If SCADA is unavailable, use monthly averages or a binary status (1.0 normal; 0.5 degraded).

- **Energy data availability (site policy).** *SCADA sites:* use real-time series. *Non-SCADA sites:* use operator bulletins/uptime reports. *Minimum:* weekly utility stability reports. *If none:* set $E_{\text{energy}} = 1.0$ and mark as an assumption.

- $I(t)$ **(intermodal synergy).** Synergy is defined between modes (e.g., bus–metro, bike–rail, park&ride) and depends on interchange distance, transfer time, and integrated payments. Normalization options:

  1. **Theoretical-maximum:** $I_{\text{norm}} = I_{\text{raw}}/I_{\text{max}}$, where I_{max} assumes the minimum observed transfer distance (e.g., 10 m) for all relevant pairs.
  2. **Empirical (min–max):** $I_{\text{norm}} = \dfrac{I_{\text{raw}} - I_{\text{min}}}{I_{\text{max}} - I_{\text{min}}}$ over a rolling window.

  Use one method consistently and document the window/assumptions.

3 Parameter JSON (excerpt for submission; full in code archive)


```
{
  "alpha_1": 0.25, "alpha_2": 0.25, "alpha_3": 0.25, "alpha_4": 0.25,
  "beta_1": 0.57, "beta_2": 0.43,
  "omega_1": 0.5, "omega_2": 0.5,
  "zeta": 0.1, "theta": 0.1, "pi": 0.1,
  "chi": 1.0,
  "kappa": 0.1,
"delta_transport": 1.0, "delta_environment_weather": 0.9,
"delta_utilities": 0.9, "delta_social": 0.9,
  "tau_days": 6,
  "omega_src_official": 1.0, "omega_src_operator": 0.9,
  "omega_src_sensor": 0.9, "omega_src_social": 0.6,
  "omega_src_survey": 0.6,
```

```
  "mu": 0.15, "lambda": 0.70
}
```

Notation: ω_1, ω_2 are infrastructure weights (energy, ICT). ω_{src} refers to source reliability weights in the disturbance mapping. These are unrelated.
Weight sets are normalized to sum to 1.0 within each category:

$$\sum \alpha = 1, \quad \sum \beta = 1, \quad \sum \omega = 1.$$

4 Core Equation

The antifragile potential $E(t)$ is an aggregate KPI $\bar{E}(x)$ per period, evolving via relaxation:

$$\dot{E} = -\kappa \left(E - \bar{E}(x)\right)$$

Note: This is a convergence model, not an equilibrium model. Future work will extend this to include traffic assignment equilibrium and modal choice equilibrium. The relaxation equation is adapted from control theory and dynamical systems, commonly used in resilience studies. From our 123-study review, we synthesized indicators that affect system performance into the target function $\bar{\text{E}}$(x). Where the target is:

$$\bar{E}(x) = \lambda \, \bar{E}^{+}(x) - (1-\lambda) \, \bar{E}^{-}(x), \qquad \lambda \in (0,1),$$

$$\text{Range:} \quad \bar{E} \in [\, \lambda - 1, \ \lambda \,]$$

$$\text{Normalisation:} \quad \bar{E}_{\text{target}} = \bar{E} - \lambda + 1$$

$$\bar{E}^{+}(x) = \alpha_1 M + \alpha_2 R + \alpha_3 A + \alpha_4 H_{\text{norm}} + \omega_1 E_{\text{energy}} + \omega_2 E_{\text{ICT}} + \zeta I + \pi P + \mu B,$$

$$\bar{E}^{-}(x) = \beta_1 D_{\text{eff}} + \beta_2 S + \theta Q.$$

$C(t)$ denotes the entropy of the origin–destination (OD) flow distribution, measuring how evenly trips are spread across OD pairs (each OD pair = one ordered origin–destination cell in the OD matrix).

$$\psi(C) = H_{\text{norm}}, \qquad H_{\text{norm}} = \frac{C}{C_{\max}(\text{current\_topology})},$$

where C is Shannon entropy (C in nats).
With $\kappa = 0.1$ (as documented in the parameter table; typical convergence in ~ 10 periods), the discrete update rule is:

$$E_{t+1} = E_t - \kappa \left(E_t - \bar{E}(x_t)\right).$$

Each city has unique core equation and equilibrium based on topology, demand patterns, and infrastructure.

Example. Consider a network with 10 zones. The number of OD pairs is $N_{\text{OD}} = 10 \times 9 = 90$. Under current topology $C_{\max}(\text{current\_topology}) = \ln(90) \approx 4.50$. Suppose the Shannon entropy of the observed OD distribution is $C = 3.5$ nats. Then

$$H_{\text{norm}} = \frac{3.5}{C_{\max}(\text{current\_topology})} = \frac{3.5}{4.5} \approx 0.78, \qquad \psi(C) = 0.78$$

Note: If network topology changes (e.g., during a disruption that closes routes), $C_{\max}$ must be recalculated. For example, if disruption reduces accessible OD pairs to 70, then $C_{\max}(\text{disrupted\_topology}) = \ln(70) \approx 4.25$, changing the normalization basis. Assume the baseline performance is $E_0 = 0.60$, the recovery target is $\bar{E} = 0.67$ (raised, in part, due to the high diversity reflected in $\psi(C) = 0.75$), and $\kappa = 0.1$. Physical Interpretation of $H_{norm} = 0.75$ means flows are well-distributed across 75% of the maximum possible entropy for the current network topology. With $\alpha_4 = 0.25$ (after normalization), this entropy level contributes $0.25 \times 0.75 = 0.1875$ to the performance target.
High entropy indicates healthy diversity when multiple viable routes are used effectively, representing "organized complexity" - like a healthy ecosystem with diverse pathways rather than over-dependence on single routes.
High entropy indicates healthy diversity when accompanied by maintained performance metrics, distinguishing organized complexity from chaotic dispersion.
The first update step is:

$$E_1 = E_0 - \kappa\left(E_0 - \bar{E}\right) = 0.60 - 0.1\,(0.60 - 0.67) = 0.60 - 0.1 \times (-0.07) = 0.607.$$

Thus the system moves 10% of the way from the current state (0.60) toward the target (0.67). Repeated iterations converge monotonically to $\bar{E}$, as $|1 - \kappa| = 0.9 < 1$ ensures stability.
In this example, with $\alpha_4 = 0.25$ and $\psi(C) = 0.75$, the entropy component contributes $\alpha_4 \times \psi(C) = 0.25 \times 0.75 = 0.1875$ to the target score $\bar{\text{E}}$(x). This significant contribution reflects how flow diversity substantially enhances system performance. Without this entropy boost, the target would be lower, demonstrating why distributed flows are crucial for antifragility.
The system's equilibrium evolves through:

- Short-term: Daily fluctuations ($\kappa = 0.2$-0.3)
- Medium-term: Seasonal patterns, policy changes
- Long-term: Infrastructure updates, behavioral shifts
- Post-disruption: Permanent improvements if AF $>$ contextual threshold

4.1 Quantifying Equilibrium Shifts

Equilibrium change is measured by:

$$\Delta E^* = E^*_{\text{post}} - E^*_{\text{pre}}$$

where E^* represents a stable state

$$\frac{dE}{dt} \approx 0.$$

Components:

- Direct performance: ΔE
- Stability improvement: $\Delta(\text{Var}(E))$
- Recovery speed: $\Delta\kappa_{\text{effective}}$
- Robustness: $\Delta\big(\min(E \text{ during next disruption})\big)$

4.2 Dynamic Entropy Normalization

During disruptions, network topology changes require recalculation of maximum entropy:

$$C_{\text{normalized}}(t) = \frac{C_{\text{actual}}(t)}{C_{\max}\big(\text{topology}_t\big)}$$

Where:

$$\begin{aligned}
\text{Pre-event:} \quad & C_{\max} = \ln\big(N_{OD_{\text{original}}}\big) \\
\text{During disruption:} \quad & C_{\max} = \ln\big(N_{OD_{\text{available}}}\big) \\
\text{Post-event with new infrastructure:} \quad & C_{\max} = \ln\big(N_{OD_{\text{new}}}\big)
\end{aligned}$$

Example:

$$\begin{aligned}
\text{Normal operations:} \quad & \text{100 OD pairs} \quad \Rightarrow \quad C_{\max} = \ln(100) = 4.605 \\
\text{During metro closure:} \quad & \text{80 accessible pairs} \quad \Rightarrow \quad C_{\max} = \ln(80) = 4.382 \\
\text{After adding bus routes:} \quad & \text{110 pairs} \quad \Rightarrow \quad C_{\max} = \ln(110) = 4.700
\end{aligned}$$

This ensures $\psi(C)$ reflects genuine flow diversity relative to current possibilities, not pre-disruption topology. Without this adjustment, entropy would be artificially deflated during disruptions or inflated after infrastructure expansion.

Implementation Validation

The code archive includes unit tests that verify:

- $C_{\max}(\text{topology}_t)$ updates correctly when OD pairs change.
- $\psi(C) \in [0, 1]$ under all network configurations.
- Entropy calculation handles edge cases (single route, complete disconnection).

4.3 Calibration of κ

The adjustment speed κ is NOT fixed but calibrated for each city based on:

- Historical recovery data (primary method)
- System type (transport: 0.2-0.3, infrastructure: 0.05-0.1)
- Event severity (major disruptions may temporarily reduce κ)

κ is recalibrated annually or after major events. Cities start with $\kappa = 0.1$ and adjust based on observed convergence rates.

5 Calibration Protocol

• Two-stage: (1) AHP expert elicitation (e.g., 20 scholars / 10 practitioners) for relative weights. (2) Elastic-net regression on 2019–2024 events to maximise out-of-sample R^2 on ΔE.
• Final weights as a convex combination: 0.6 expert + 0.4 data (document ranges and cross-validated λ).
• Implementation: Python; dynamics via a discrete update; archive code and seeds.

Data Availability Tiers: - Full (>70% coverage): All indicators active - Partial (30-70%): Core indicators + proxies - Minimal (<30%): M(t), S(t),D_{eff} only + estimates
Correlation Matrix Projection: If pilot data yields a non-positive-semidefinite correlation matrix:

1. Apply Higham (2002) nearest correlation matrix projection.
2. Preserve diagonal elements (all = 1.0).
3. Minimize Frobenius norm of correction.
4. Verify all eigenvalues ≥ 0 in the corrected matrix.
5. Document any significant changes from theoretical expectations.

Correction Documentation
For reproducibility, the archive includes:

- Pre-correction eigenvalues and condition number.
- Post-correction eigenvalues (all ≥ 0).
- Frobenius norm of correction matrix $\|R_{\text{corrected}} - R_{\text{original}}\|_F$.
- Flag if any correlation changes by > 0.15:
  `"Large correction applied to R-C correlation: 0.78 → 0.65"`.

This ensures transparency about any significant modifications to the theoretical correlation structure.

6 Sensitivity Analysis (Sobol)

Global sensitivity on mean E and settling time (time to $|E - \bar{E}| \leq 5\%$ of the initial gap).
First-order S_i measures individual impact; total-order S_{T_i} includes interactions.
Example placeholder results (replace with project-run outputs):

Metric	**Top Parameters**	**Values**
First-Order (mean E)	K_R, R, α_4, E_{ICT}, α_3	0.1260, 0.1223, 0.0759, 0.0371, 0.0307
Total-Order (mean E)	K_R, R, α_4, E_{ICT}, α_3	0.1614, 0.1573, 0.1035, 0.0564, 0.0482

How First-Order and Total-Order Indices are determined?
First-Order (Si):
• Fix parameter i at different values • Let all other parameters vary randomly • Calculate variance of output when i is fixed • Si = (Total variance - Variance when i fixed) / Total variance
Total-Order (STi):
• Let all parameters except i vary • Calculate variance from i and its interactions • STi = Variance due to i (including interactions) / Total variance

6.1 Indicator Correlations

Positive Correlations:

- $R \uparrow \Rightarrow C \uparrow$ (0.65): More redundancy enables flow diversity
- $I \uparrow \Rightarrow M \uparrow$ (0.45): Better integration increases throughput
- $P \uparrow \Rightarrow A \uparrow$ (0.55): Pricing drives adaptation

Negative Correlations:

- $D_{\text{eff}} \uparrow \Rightarrow M \downarrow$ (-0.70): Disruptions reduce throughput
- $S \uparrow \Rightarrow C \downarrow$ (-0.40): Stress concentrates flows
- $P \uparrow \Rightarrow Q \uparrow$ (-0.35): Pricing may hurt equity

Trade-offs requiring balance:

- Efficiency (M) vs. Resilience (R)
- Performance (E) vs. Equity (Q)
- Speed (S low) vs. Distribution (C high)

6.2 Correlation Matrix Validation

The complete correlation matrix must be constructed and validated:

	M	R	A	C	D_{eff}	S	Q	P	I
M	1.0	0.3	0.4	0.2	-0.70	-0.5	-0.1	0.2	0.45
R	0.3	1.0	0.2	0.65	-0.4	-0.6	0.0	0.1	0.3
A	0.4	0.2	1.0	0.3	-0.3	-0.2	0.1	0.55	0.2
C	0.2	0.65	0.3	1.0	-0.2	-0.4	0.1	0.3	0.4
D_{eff}	-0.70	-0.4	-0.3	-0.2	1.0	0.7	0.2	-0.1	-0.3
S	-0.5	-0.6	-0.2	-0.4	0.7	1.0	0.3	-0.2	-0.4
Q	-0.1	0.0	0.1	0.1	0.2	0.3	1.0	-0.35	0.0
P	0.2	0.1	0.55	0.3	-0.1	-0.2	-0.35	1.0	0.1
I	0.45	0.3	0.2	0.4	-0.3	-0.4	0.0	0.1	1.0

Matrix Properties Check:

- **Eigenvalues:** [computed from actual pilot data]
- **Condition number:** [must be < 1000 for numerical stability]
- **Positive semidefinite:** [verify all eigenvalues ≥ 0]

6.3 Multicollinearity Assessment

High correlations ($|r| > 0.8$) requiring attention:

- None detected in current specification
- Monitor R–C correlation (0.65) during calibration
- If $|r| > 0.8$ emerges, consider indicator consolidation

Variance Inflation Factors (VIF):

$$\text{VIF} = \frac{1}{1 - R^2} \quad \text{for each indicator}$$

- VIF > 5 indicates problematic multicollinearity.

6.4 Theoretical Sign Consistency

Expected relationships based on transport theory:
Validated signs:

✓ $S \uparrow \Rightarrow C \downarrow$ (congestion concentrates flows – Sheffi, 1985)

✓ $D_{\text{eff}} \uparrow \Rightarrow M \downarrow$ (disruptions reduce throughput)

✓ $R \uparrow \Rightarrow C \uparrow$ (redundancy enables flow diversity)

Questionable relationships requiring empirical validation:

? $P \uparrow \Rightarrow Q \uparrow$ (pricing effects on equity vary by implementation)

? $I \uparrow \Rightarrow M \uparrow$ (synergy–throughput relationship needs verification)

7 Back-Testing

7.1 Scope and Datasets

We implement two complementary back-testing tracks until Phase 2 pilot data are available:

1. **Benchmark-network experiments (no real data):** run the full pipeline on canonical networks (e.g., Sioux Falls) under stylised shocks (capacity drops, demand surges) to verify that UE/SO-consistent interventions lead to statistically significant antifragility (AF) detections.

2. **Synthetic back-tests:** generate demand/supply series with known ground-truth improvements; estimate $(\kappa, \sigma, \mathbf{w})$ on a training segment and validate out-of-sample detection power and false-positive control.

Phase 2 will add: *City A / Event Type 1* for calibration and *City B / Event Type 2* for external validation (holdout).

7.2 Experimental Design

Calibrate → Validate. Calibrate parameters on Track T1/T2 or City A, then validate on a disjoint case (Track T1/T2 held-out scenarios or City B). Report predictions for AF and errors against observed AF decisions.

Ablation. Run a minimal ablation to quantify contribution to AF:

1. drop R (redundancy) only,
2. drop C (diversity/entropy) only,
3. drop both R and C.

Report the change in effect size and decision outcomes.

7.3 Statistical Validation Protocol

Let E_{base} be the seasonally-adjusted baseline mean over a pre-event window. Define the post-event 30-day mean

$$\mu_{30} \;=\; \frac{1}{30}\sum_{t=31}^{60} E(t), \quad R \;=\; \frac{\mu_{30}}{E_{\text{base}}}.$$

We use two thresholds and decide by the maximum:

$$AF_{\text{req}} = 1 + 0.2\,(1 - E_{\text{base}}), \qquad AF_{\text{crit}} = 1 + \delta, \quad \delta = k\,\hat{\sigma}_{\text{resid}},$$

where $\hat{\sigma}_{\text{resid}}$ is the standard deviation of pre-event residuals after seasonal adjustment (STL or X-13), and k sets confidence (default $k = 1.64$ for 95% one-sided). We declare AF if

$$\textbf{Decision:} \quad \text{AF if} \quad \text{LCL}_{0.95}(R) \;\geq\; \max\{AF_{\text{req}}, AF_{\text{crit}}\}.$$

Confidence interval for the ratio. Compute a one-sided 95% lower confidence limit $\text{LCL}_{0.95}(R)$ using either:

1. *Bootstrap (recommended):* block bootstrap daily values in both windows and form the percentile lower bound of R (preserves autocorrelation); or
2. *Delta method:* approximate $\text{Var}(R)$ from window variances and apply a one-sided normal CI when independence and homoscedasticity are plausible.

10-day stability test. Require convergence within noise in the final 10 days of the 30-day window:

$$|E(t) - \mu_{30}| \;<\; 2\,\hat{\sigma}_{\text{resid}} \quad \text{for all } t \in [51, 60].$$

As a robustness check, report the maximum absolute CUSUM over $t \in [31, 60]$ and confirm it lies below the 95% control bound.

7.4 Seasonality, Lags, and Alignment

Apply consistent temporal alignment before computing $\bar{E}(x_t)$: daily aggregation for mixed-frequency inputs; carry-forward for weekly/monthly policy variables; and explicit lag structure where theoretically required. Perform seasonal adjustment on $E(t)$ and inputs prior to baseline estimation.

7.5 Sensitivity and Uncertainty

Report: (i) Sobol or Morris sensitivity of the AF decision to $(\kappa, \sigma, \mathbf{w})$ and missing-data mechanisms; (ii) ablation deltas; (iii) coverage-weighted uncertainty where proxies/imputation are used.

7.6 Reporting

For each case (Track T1/T2 or City B), report:

1. E_{base}, μ_{30}, R, $\text{LCL}_{0.95}(R)$, $\max\{AF_{\text{req}}, AF_{\text{crit}}\}$, and the AF decision;
2. stability test outcome and CUSUM summary;
3. ablation results (decision flips and effect-size changes);
4. run-time and data coverage.

8 Risks & Limitations

• Data sparsity: coverage $<70\%$ degrades accuracy; apply imputation and report uncertainty.
• Policy inertia: demand-management (P) changes may lag events; document policy timelines.
• External validity: calibrated primarily on mid-size European cities; caution on transfer.
• Model scope: extreme 'black swan' cascades are not explicitly modelled.

8.1 Methodology for Pre/Post-Event Comparison

When structural changes occur (infrastructure additions/removals), direct comparison of absolute metrics becomes invalid. The model addresses this through: Functional Equivalence Approach: Instead of comparing raw indicator values, we normalize by system capacity and demand: Comparable metrics:

$$\text{Comparable}_R = \frac{\text{Redundant Capacity} \times \text{Utilization}}{\text{Total Demand}}$$

$$\text{Comparable}_M = \frac{\text{Throughput}}{\text{Population} \times \text{Available Routes}}$$

$$\text{Comparable}_C = \frac{C_{\text{actual}}}{C_{\text{max}}(\text{current topology})}$$

Example - Redundancy Functional Comparison:

- Pre-event: $R = 0.20$ (60 spare routes, 300 total routes), demand = 1000 trips
- Post-event: $R = 0.25$ (50 spare routes, 200 total routes), demand = 900 trips

Raw comparison suggests R improved ($0.20 \rightarrow 0.25$).

Functional comparison using spare capacity per unit demand:

$$\text{Pre: Spare\_capacity\_ratio} = \frac{60\_Spare\_roots}{1000\_trips} = 0.060 \text{ spare routes per trip}$$

$$\text{Post: Spare\_capacity\_ratio} = \frac{50\_Spare\_roots}{900\_trips} = 0.056 \text{ spare routes per trip}$$

Result: Despite higher R ratio, functional spare capacity decreased by 7% $\frac{(0.056-0.060)}{0.060}$ -0.07 This shows the network is actually less robust per unit of demand served.

Example - Metro expansion changes reference:

Pre-expansion baseline: E = 0.60 with 100 OD pairs
Post-expansion: E = 0.66 with 120 OD pairs
Raw comparison: 0.66/0.60 = 1.10 Normalized comparison:

- Performance per OD: (0.66/120)/(0.60/100) = 0.92
- Functional improvement: Service quality × 1.10, Coverage × 1.20
- True AF considering both: 1.10 × 0.92 = 1.01 (marginal improvement)

9 Reproducibility & Compliance Appendix

(i) GDPR Note: perform a DPIA; apply k-anonymity (k=5) and, where appropriate, differential privacy ($\varepsilon \approx 0.1$); no raw personal data shared.
(ii) Data Management Summary: follow FAIR principles; provide anonymised sample datasets via EU-compliant portals.
(iii) Code Archive: Python script for dynamics and sensitivity (NumPy/Scipy); sample dataset; correlation matrix validation logs with pre/post eigenvalues; README with run instructions (e.g., run main.py with params.json). Archive DOI to be added..

9.1 Expected Pilot Data Availability

High Availability (>90%):

- Smart card/ticketing data (M indicator)
- Traffic counts from sensors (S indicator)
- GTFS transit schedules
- OpenStreetMap network topology

Medium Availability (50–70%):

- Household travel surveys (annual/biannual)

- Mobile phone CDR data (privacy restrictions apply)
- Parking occupancy
- Weather data

Low Availability (<30%):

- EV charging grid data (varies by utility cooperation)
- Real-time ICT system status
- Detailed equity data by demographic

Pilot-Specific Adaptations:
Cities with <70% data coverage will use proxy indicators and confidence intervals as documented in Section 8.

10 Implementation Pseudocode

For each time step t:

1. Read inputs
$$x(t) = [M, R, A, C, D_{\text{eff}}, S, Q, P, E_{\text{energy}}, E_{\text{ICT}}, I, B].$$

2. Compute positive and negative blocks
$$\begin{aligned}\bar{E}^{+} &= \alpha_1 M + \alpha_2 R + \alpha_3 A + \alpha_4 H_{\text{norm}} \\ &\quad + \omega_1 E_{\text{energy}} + \omega_2 E_{\text{ICT}} + \zeta I + \pi P + \mu B, \\ \bar{E}^{-} &= \beta_1 D_{\text{eff}} + \beta_2 S + \theta Q, \\ H_{\text{norm}} &= \frac{C}{C_{\max}(\text{current\_topology})}.\end{aligned}$$

3. Compute bounded target
$$\bar{E} = \lambda\, \bar{E}^{+} - (1-\lambda)\, \bar{E}^{-}$$

4. Map to [0,1] range
$$E_{\text{target, normalized}} = E - \lambda + 1 \qquad \text{// maps } [\lambda - 1,\, \lambda] \ \rightarrow\ [0,1]$$

5. Update performance
$$E \ \leftarrow\ E - \kappa\,(E - \bar{E}_{\text{bounded}}).$$

6. Log E, $\bar{E}_{\text{bounded}}$, and inputs for plots and checks.

11 Convergence Verification Criteria

Improved system performance specifically means:

- Travel time reduction (>10%)
- Mode share optimization (closer to sustainable targets)
- Reduced vulnerability (lower peak S during events)
- Better equity (lower Q)
- All sustained for 30+ days without external intervention

The model tracks a performance score E(t) for the urban mobility system. D2.1 produced a catalogue of events (domain, scale, severity, time, source, AOI). We convert those events into a composite disturbance D_eff(t) that, together with stress S(t) on affected areas, shifts the target score $\bar{E}(x)$. E(t) then moves gradually toward $\bar{E}(x)$ at speed κ, like a thermostat approaching a set-point.
The current model uses a convergence approach toward a target state:

$$E^* \quad \text{where} \quad \frac{dE}{dt} = 0 \quad \text{and} \quad \frac{d^2E}{dt^2} = 0$$

which differs from true equilibrium models such as:

- **Pareto equilibrium:** No improvement possible without trade-offs.
- **Nash equilibrium:** No agent benefits from unilateral change.
- **Wardrop equilibrium:** Equal travel costs on used routes.

The equilibrium version implements true multi-layer equilibria, including Wardrop (traffic) and discrete choice (modal) equilibrium. The current model focuses on **performance convergence** for computational tractability.

11.1 Event→Model linkage (core)

For an event e with domain d, severity level $L \in \{1, \dots, 5\}$, start time t_e, and source type src, define

$$D_{\text{eff}}(e,t) = \omega_{\text{src}}(\text{src}) \cdot \exp\!\Big(-\tfrac{t-t_e}{\tau}\Big) \cdot \delta_d \cdot \frac{L-1}{4}.$$

Aggregate across concurrent events:

$$D_{\text{eff}}(t) = \sum_e D_{\text{eff}}(e,t).$$

Here $D_{\text{eff}}(t)$ replaces raw disturbance terms and co-varies with $S(t)$ computed on links within the event's area(s) of interest (AOIs).

Scale and Magnitude Guidelines

- Individual events: $D_{\text{eff}}(e,t) \in [0,1]$ (severity $L = 5$, immediate, reliable source).

- Typical magnitudes:
  - Minor ($L = 1$–2): 0.1–0.3,
  - Major ($L = 4$–5): 0.5–0.8.
- Multiple concurrent events:

$$D_{\text{eff}}(t) = \sum D_{\text{eff}}(e, t)$$

  with soft cap at 1.0:

$$D_{\text{eff,final}} = \min\bigl(1.0,\ D_{\text{eff,raw}}\bigr) \quad \text{or use softplus: } \ln\bigl(1 + e^{D_{\text{eff,raw}}}\bigr).$$

Calibration Check:
With $\beta_1 \approx 0.57$, a major event ($D_{\text{eff}} = 0.6$) reduces target by 0.34 points.
Users should verify this penalty magnitude aligns with expected system response. Scale and Magnitude Guidelines

- Individual events: $D_{\text{eff}}(e, t) \in [0, 1]$ (severity $L = 5$, immediate, reliable source).

- Typical magnitudes:
  - Minor ($L = 1$–2): 0.1–0.3,
  - Major ($L = 4$–5): 0.5–0.8.
- Multiple concurrent events:

$$D_{\text{eff}}(t) = \sum D_{\text{eff}}(e, t)$$

  with soft cap at 1.0:

$$D_{\text{eff,final}} = \min\bigl(1.0,\ D_{\text{eff,raw}}\bigr) \quad \text{or use softplus: } \ln\bigl(1 + e^{D_{\text{eff,raw}}}\bigr).$$

11.2 Equilibrium Verification Criteria

New equilibrium confirmed when:

- Convergence: |E(t) - E(t-1)| $<$ 0.01 for 5 consecutive periods
- No drift: Linear regression slope of E(t) $<$ 0.001
-
$$\forall x_i : \ \bigl|x_i(t) - \text{mean}(x_i)\bigr| < 0.05 \times \bigl(\max(x_i) - \min(x_i)\bigr),$$
$$\text{where } \max(x_i), \min(x_i) \text{ are computed over the 30-day stability window.}$$
- No active interventions required to maintain performance

11.3 Normalisation

Indicators are normalised over a rolling 30-day window using

$$\hat{x} = \frac{x - x_{\min}}{x_{\max} - x_{\min}} \in [0, 1].$$

The baseline x_0 is tracked separately for antifragility (AF) comparisons.
Example. In the last 30 days, trips per capita per hour ranged 0.8–1.6. Today $x = 1.3$. Then

$$\hat{x} = \frac{1.3 - 0.8}{1.6 - 0.8} = \frac{0.5}{0.8} = 0.625.$$

Note: always use $x_{\min}$ from the 30-day window, not day-1 x_0 (which is tracked separately for AF).

11.4 Variable dictionary - what to send & examples

Code	**Plain-language name**	**What it means / how to compute**	**Unit**	**Frequency**	**Typical source**	**Simple example**
$M(t)$	Mobility throughput	Average multi-modal trips per person per hour (sum trips across modes ÷ population ÷ hours).	trips h^{-1} cap^{-1}	Daily average	Smart-card, CDR, GPS probes	2.4M trips/person/day with pop 1.0M → 2.4 trips/person/day ≈ 0.10 trips/person/hour.
$R(t)$	Redundancy ratio	Dormant links ÷ active links in pre-event snapshot (links that can carry flow but are idle vs. carrying).	–	Weekly	OSM topology + counts, sensors	300 active, 60 dormant → R = 0.20.
$A(t)$	Adaptation velocity	Rate of change of modal-share diversity over time; finite difference of a diversity index per day.	$\% \cdot \text{day}^{-1}$	Bi-weekly	Household/intercept surveys	Diversity 0.62 vs 0.58 over 14 days → (0.62−0.58)/14 ≈ 0.0029 day^{-1}.
$C(t)$	Network entropy of OD matrix	Shannon entropy of OD flows, normalised by maximum entropy of current available topology; higher = flows are more spread-out.	nats (normalised)	Daily	CDR, Bluetooth, APC/AVL	If flows are evenly distributed over 4 ODs, H is high; if 90% on one OD, H is low.
$D_{\text{eff}}(t)$	Event disturbance	Domain/recency/source-weighted disturbance index.	–	Hourly → daily aggregate	D2.1 event catalogue + operator/official logs	Transport event, $L = 4$, 12h old, operator source, τ = 6d, $\delta = 1.0$, $\omega_{\text{src}} = 0.9$ →≈ 0.621.
$S(t)$	Stress index	Σ over links in AOIs of (vol/cap) × speed_drop; aggregates congestion and capacity strain.	–	Hourly → daily aggregate	Loop detectors, floating-car data	If vol/cap= 0.9 and speed_drop= 0.2 on a link → 0.18; sum across links.
$Q(t)$	Equity penalty	Stratified Theil-T of accessibility minutes across income or neighbourhood groups.	–	Monthly	Census + GTFS travel times	If low-income groups face 15% longer mean access times, Q increases.
$P(t)$	Demand-management	Composite index in [0,1] from time-varying fares via min–max by mode/time, aggregated with ridership weights.	index (0–1)	Daily	Operator APIs, open fares	Peak surcharge +20% → index rises (e.g., 0.65→0.72).
$B(t)$	Behavioral compliance	Weighted average of: route guidance compliance (% following apps), mode shift willingness (survey), temporal flexibility (% able to shift travel time)	index (0–1)	Weekly	App analytics, surveys, smart card patterns	0.7 = 70 follow guidance, moderate flexibility
$E_{\text{energy}}(t)$	EV-grid resilience	Composite of uptime × spare capacity for EV charging during events.	–	Daily	Grid SCADA, DSO reports	Uptime 99.5% × capacity index 0.8 →≈ 0.796.
$E_{\text{ICT}}(t)$	ICT uptime (during event)	Percentage uptime of telecom/-control systems critical to operations.	%	Event-specific	Operator logs / incident reports	30 minutes downtime in a 10-hour event → 95% uptime.
$I(t)$	Cross-layer synergy	Graph-based proximity between modal layers (e.g., bike–rail hubs); sum of inverted transfer distances.	–	Weekly	Multilayer transport graph	If new hubs cut average transfer distance by 20%, I increases.

⋆ Regarding C(t), disruptions can either:

• Concentrate flows (routes blocked → everyone on same detour) → Low C
The key is whether spread flows are efficient (good) or chaotic (bad). We validate this by checking if performance E also improves.
• Spread flows (people find various alternatives) → High C
⋆ In $E_{energy(t)}$, during events, measures both: *Uptime*: % of chargers operational (e.g., 95%), *Spare capacity*: Unused charging capacity (e.g., 60% available), *Composite*: $0.95 \times 0.60 = 0.57$ This captures if EVs can still charge during disruptions.

11.5 Parameter list — meanings and handling

Parameter	Meaning / role	Typical source / note
$\alpha_1 \ldots \alpha_4$	Weights on M, R, A, ψ(C)	Elicited via AHP; refined via elastic-net on 2019–2024 events.
$\beta_1,\ \beta_2$	Weights on D_eff and S (penalties)	Expect β>0; calibrated against outcomes.
$\omega_1,\ \omega_2$	Weights on energy and ICT resilience	Higher values increase target performance under reliable infra.
μ	Weight on behavioral compliance B	Higher μ means user cooperation more critical for performance
ζ	Weight on synergy I	Captures benefit of cross-modal integration.
θ	Weight on equity penalty Q	Higher θ penalises inequitable access.
π	Weight on demand management P	Reflects effectiveness of pricing/management levers.
χ	I(t) normalization: Min-max over 6-month rolling window or theoretical maximum	Tune to keep I within [0,1] after normalisation.
κ	Adjustment speed of E toward Ē	κ=0.1 implies ∼10 periods to close most of the gap.
δ (domain weights)	Relative impact of domains on D_eff	Init from D2.1 prevalence & impact; refine with elastic-net.
τ (recency half-life)	Decay of older events: exp(−Δt/τ)	Choose so past week dominates; cross-validate.
ω_src (source reliability)	Weight by source: official/operator > sensor > social/survey	Set tiers during pilot; combine with ESS-based shrinkage.
λ	mixes positive/negative blocks in $\bar{E} = \lambda\, \bar{E}^{+} - (1-\lambda)\, \bar{E}^{-}$	expert elicitation (AHP) or regression to historical targets; fixed per city until next recalibration.
Constraint	$\Sigma(\alpha_1, \alpha_2, \alpha_3, \alpha_4) = 1.0$, $\Sigma(\beta_1, \beta_2) = 1.0$, $\Sigma(\omega_1, \omega_2) = 1.0$	Ensures Ē(x) remains bounded

$\star$ Start with μ = 0.15 (moderate importance). Cities with high social capital might use μ = 0.20; low-trust environments μ = 0.10.

11.6 Spatial & temporal resolutions (core)

Spatial: 1 km^2 grid aggregation for urban centres and mixed-use areas; AOIs snapped to grid and network links.
Temporal: hourly updates for transport-facing indicators; baselines computed on rolling 6-month windows.
Baseline Comparability: For valid AF calculation, pre-event baseline metrics are stored both as absolute values and normalized functional capacities. Post-event comparison uses the normalization appropriate to the structural change type (temporary vs permanent).
Measurement Schedule:

- Baseline: 30 days pre-event average
- During event: Daily measurements
- Recovery: Daily for 14 days
- Validation: Weekly for 4 weeks
- Long-term: Monthly for 6 months
- Annual assessment to confirm sustained AF > threshold

Temporal Alignment Protocol
For daily $\bar{E}(x_t)$ computation with mixed-frequency indicators:

- Hourly indicators (S): use most recent hourly value.
- Daily indicators (M, C): use current day's value.
- Weekly indicators (B): last observation carried forward (LOCF).
- Monthly indicators (Q): LOCF with linear interpolation if $>$ 45 days old.

Lag Structure:

- Real-time indicators: no lag (S, M, C).
- Survey-based indicators: 1-period lag to account for collection delay (B, Q).
- Event indicators (D_{eff}): exponential decay from event time.

12 Antifragility Threshold and Model Linearity

The relaxation step E_{t+1} = E_t − κ (E_t − Ē(x_t)) is linear in E, but the inputs x_t are nonlinear functions of data: D_{eff} uses exponential recency and weights; ψ(C) depends on normalised entropy; S aggregates link-level nonlinearities. Breaking down the nonlinearities:

- **Exponential recency:**

  $$\exp\left(-\frac{\Delta t}{\tau}\right)$$

  makes recent events matter more (e.g., 12h old ⇒ 0.92 impact, 48h old ⇒ 0.71).

- **Normalized entropy:**

$$\psi(C) = H_{\text{norm}} = \frac{C}{C_{\max}(\text{current\_topology})}$$

  creates a nonlinear boost from flow diversity.

- **Link nonlinearities:**

$$S = \sum \left(\frac{\text{volume}}{\text{capacity}} \right) \times \text{speed}_{\text{drop}}$$

  captures congestion's exponential growth.

Hence, the overall input–output mapping is nonlinear even though the convergence mechanism is linear (stable for $0<\kappa<2$, ensuring convergence).

12.1 Statistical AF Threshold Determination

Methodology:

1. Baseline noise estimation: calculate σ_{residual} from 6-month pre-event $E(t)$ data.
2. Seasonal modeling: fit ARIMA(p, d, q) to remove systematic patterns.
3. Statistical threshold: $\delta = k \times \sigma_{\text{residual}}$, where $k = 2.58$ (99% confidence).
4. City-specific threshold: $AF_{\text{critical}} = 1 + \delta$.

Hypothesis Test:

$$H_0 : \mu_{30a} \leq (1 + \delta)E_{\text{base}} \qquad \text{(no antifragile improvement)}$$

$$H_1 : \mu_{30a} > (1 + \delta)E_{\text{base}} \qquad \text{(statistically significant antifragility)}$$

Implementation:

- Calculate 30-day post-recovery mean: μ_{30a}.
- Compute one-sided 95% confidence interval.
- Declare antifragility if $CI_{\text{lower}} > (1 + \delta)E_{\text{base}}$.

Example Applications:

- High-noise city: $\sigma_{\text{residual}} = 0.05 \Rightarrow AF > 1.129$ required.
- Low-noise city: $\sigma_{\text{residual}} = 0.02 \Rightarrow AF > 1.052$ required.

Advantages:

- Controls false discovery rate at specified α level.
- Accounts for city-specific measurement characteristics.
- Provides statistical confidence in antifragility claims.
- Maintains comparability through standardized effect sizes.

13 Future Work

This convergence model will be extended to a full equilibrium framework incorporating Wardrop traffic assignment and discrete choice modal equilibrium in Phase 2 of the project.

Appendix