Infrastructure Transitions Precede Congestion: A Regime-Transition Framework for UAV Logistics and Airspace Planning

Bingqing Liu , Qunting Yang * , Chunsheng Xie , Yongxuan Song , Zhang Wen , Qian Yang , Pengyu Jin

Posted Date: 1 June 2026

doi: 10.20944/preprints202605.2095.v1

Keywords: UAV logistics; airspace capacity; depot location; regime transition; critical demand threshold; two-phase mechanism; infrastructure planning

 doi:10.20944/preprints202605.2095.v1

Article

Infrastructure Transitions Precede Congestion: A Regime-Transition Framework for UAV Logistics and Airspace Planning

Bingqing Liu [1,2], Qunting Yang [1,*], Chunsheng Xie [1], Yongxuan Song [3], Zhang Wen [1], Qian Yang [1] and Pengyu Jin [1]

[1] College of Air Traffic Management, Civil Aviation University of China, Tianjin 300300, China
[2] North China Municipal Engineering Design & Research Institute Co., Ltd., Tianjin, China
[3] School of Graduate Studies, Lingnan University, Hong Kong, China
* Correspondence: qtyang@cauc.edu.cn

Abstract

Urban UAV logistics planning optimizes depot locations for delivery distance while treating airspace capacity as a fixed constraint. We show that this decoupling is not merely a modeling simplification — it produces a structural error that cost monitoring cannot detect in time. When depot configurations begin diverging under airspace capacity pressure, the window for low-cost infrastructure correction closes before any cost signal appears. The central finding is a lead–lag relationship between structural and cost responses: at demand levels where total cost differences remain negligible, depot layouts already diverge substantially under airspace capacity pressure. Cost is a lagging indicator, not a planning trigger. To close this gap, we derive the Ω index, a closed-form demand threshold with three city-level parameters — capacity gap ΔC (from building-footprint data via a FAR volumetric method), depot distance differential Δd (from road-network geometry), and mean OD delivery distance d (from weighted building-centroid sampling) — separating a regime where distance-based planning is analytically sufficient from one where airspace capacity must be internalized into the location decision. Applied to Dongli District, Tianjin, the threshold lies above current UAV density but within a foreseeable planning horizon given observed growth rates in Chinese pilot cities. The depot relocation distance D_{reloc} provides a leading structural indicator, detectable well before any cost signal emerges. The Ω index gives planners a computable early-warning trigger for airspace-integrated infrastructure review — before congestion, not after.

Keywords: UAV logistics; airspace capacity; depot location; regime transition; critical demand threshold; two-phase mechanism; infrastructure planning

1. Introduction

1.1. The Problem: A Hidden System Transition

China's low-altitude airspace is undergoing a historic regulatory transformation that makes the research question of this paper both timely and practically urgent. The revised Civil Aviation Law, which took effect 1 July 2026, establishes a two-tier classification: W-class airspace (0–120 m, simple registration via the national UOM platform) and G-class (120–300 m, streamlined approval). Critically, all three operational layers modeled in this paper (L0: 0–40 m, L1: 40–80 m, L2: 80–120 m) fall entirely within W-class: the commercially open segment where UAV logistics will scale first and fastest. China's unified UOM flight management platform makes corridor-level demand flow f directly observable from mandatory registration records, enabling real-time monitoring of the Ω index that is not yet feasible under UTM architectures elsewhere. Six pilot cities (Hefei, Hangzhou, Shenzhen, Suzhou, Chengdu, and Chongqing) have received delegated authority over airspace below 600 m, making city-level planning tools such as the Ω framework immediately actionable by local governments.

 doi:10.20944/preprints202605.2095.v1

Urban UAV logistics is approaching an inflection point. As cities transition from pilot projects to early-scale commercialization (Otto and others, 2018; Rimjha and others, 2023; Stolaroff and others, 2018), a structural tension is emerging between two dimensions of infrastructure planning that have historically been treated as independent. Logistics optimization, rooted in the location-routing problem (LRP), minimizes delivery distance. Urban airspace management (UTM) allocates capacity in shared low-altitude corridors. These paradigms produce different answers about where depots should be sited. At low demand, the gap between their answers is inconsequential. At high demand, it is not. The question this paper addresses is: **at what point does ignoring this gap become a structural error?** Figure 1 illustrates the four-stage progression: Panel A (unconstrained, distance-optimal siting), Panel B (structural adaptation under airspace restrictions), Panel C (regime transition at λ^*), and Panel D (congestion-dominated reconfiguration).

The answer cannot be read from aggregate cost metrics. When demand is low, a logistics planner monitoring total system cost sees no signal that airspace constraints matter: plans accounting for airspace capacity and plans that ignore it produce nearly identical costs (within 2.8% in our experiments), even though depot configurations under the two approaches have already diverged substantially (mean relocation distances exceeding 1,900 m, customer reassignment rates above 90%). Cost is a lagging indicator.

Infrastructure transitions precede congestion. Depot configurations begin shifting in response to airspace constraints at demand levels where cost metrics show no anomaly. Monitoring total cost alone will fail to detect this early phase of system transition.

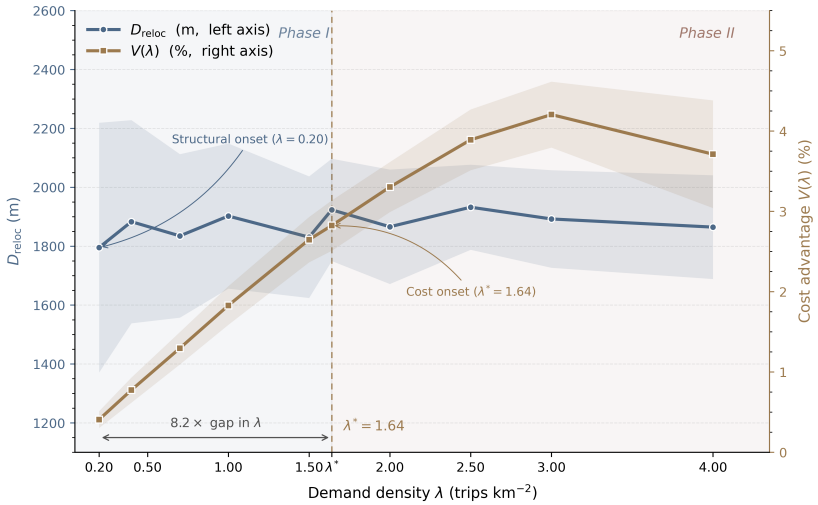

Figure 1. Regime-transition mechanism in UAV logistics infrastructure planning. Four panels trace the progression from unconstrained distance-optimal siting (Panel A) through structural adaptation (Panel B) and regime transition at $\lambda^* = 1.64$ trips/km^2 (Panel C) to congestion-dominated reconfiguration (Panel D). Cost monitoring signals a problem only in Panel D; the infrastructure transition is already underway in Panels B–C.

1.2. Research Gaps: Three Missing Pieces

Despite extensive parallel literatures on UAV facility location (Dorling et al., 2017; Macrina and others, 2020; Murray and Chu, 2015) and urban airspace management (Kopardekar and others, 2016; Levin et al., 2024; Weng and others, 2025), three critical gaps remain.

Gap 1: No structural transition theory for UAV logistics. Classical transportation theory demonstrates that capacity-constrained systems can exhibit nonlinear phase transitions as demand increases (Daganzo, 1994; Vickrey, 1969). Whether analogous mechanisms exist in UAV logistics is unknown.

Gap 2: No decision criterion for when to integrate the two paradigms. LRP studies optimize depot locations assuming unconstrained airspace (Kitjacharoenchai and others, 2019; Murray and Chu,

 doi:10.20944/preprints202605.2095.v1

2015). UTM studies model corridor congestion with depot locations fixed as exogenous inputs (Levin et al., 2024; Weng and others, 2025). No framework determines when airspace capacity should be internalized into the location decision.

Gap 3: No early-warning indicator for infrastructure planning. Even if a planner understands that a transition exists, conventional metrics (total cost, congestion cost) will not detect it in time.

1.3. This Paper: A Regime-Transition Framework

We address all three gaps with a unified analytical framework built around one central insight: the phase transition in UAV logistics infrastructure is structurally analogous to congestion phase transitions in road networks, but its onset precedes any observable congestion.

Contribution 1 (Structural): We identify and formally characterize a two-phase regime transition in UAV logistics systems, with an analytically derived boundary λ^* separating a structural sensitivity region (Phase I) from a cost escalation region (Phase II).

Contribution 2 (Analytical): We derive the Ω index, a closed-form demand threshold with three city-level input parameters (capacity gap ΔC, depot distance differential Δd, and mean OD distance d). $\Omega > 1$ defines the condition under which airspace-aware depot siting is structurally superior to distance-minimization.

Contribution 3 (Empirical): We demonstrate that depot relocation distance D_{reloc} is a sensitive early-warning indicator of the approaching transition, detectable at $\lambda = 0.20$ trips/km^2—the lowest demand level tested—while the cost signal first becomes positive at $\lambda^* = 1.64$, a gap of 8.2$\times$ in demand density (1.64/0.20).

Table 1 positions this paper relative to existing literature.

Table 1. Marginal Contribution Relative to Existing Paradigms.

Dimension	Existing LRP	Existing UTM*	This Paper
Depot location	Optimized	Fixed (exogenous)	Optimized
Airspace capacity	Ignored	Modeled	Endogenous + threshold
Decision criterion	None	None	Ω index (closed-form)
Phase transition	None	None	Two-phase mechanism
Early-warning indicator	None	None	D_{reloc} (leading)
Cost vs. structure	None	None	Lead–lag relationship formalized

Note: LRP = Location-Routing Problem; UTM = Urban Traffic Management (airspace); *UTM column covers existing airspace management models in which depot/vertiport locations are treated as fixed exogenous inputs; D_{reloc} = mean depot relocation distance between airspace-aware and distance-only plans; "leading" indicates that D_{reloc} detects the structural transition before cost metrics escalate.

2. Literature Review

Two parallel literatures bear directly on this paper, but neither addresses their intersection. The UAV location-routing problem (UAV-LRP) has expanded substantially since Murray and Chu (2015) introduced the flying sidekick TSP and Dorling et al. (2017) formalized drone delivery routing, with subsequent extensions to multi-depot configurations (Chung et al., 2020), heterogeneous fleets, and risk-aware objectives. The UTM literature has developed sophisticated capacity models: Weng and others (2025) demonstrate BPR-consistent congestion in low-altitude corridors; Levin et al. (2024) introduce the first UAV network design model with BPR-style corridor congestion; Cummings and Mahmassani (2024) derive airspace flow relations for advanced air mobility; and Mercan et al. (2025) and Revillod et al. (2025) address vertiport location and sizing. What unifies both literatures is a shared structural decoupling: UAV-LRP studies assume unconstrained airspace, while UTM studies fix depot locations as exogenous inputs. No existing formulation jointly endogenizes airspace capacity into the depot location decision.

The structural motivation for a regime-transition approach comes from classical transportation science. Daganzo (1994)'s cell transmission model and Vickrey (1969)'s analysis of infrastructure

 doi:10.20944/preprints202605.2095.v1

investment timing both establish that the structural onset of a congestion regime and its observable cost consequences are temporally separated—a phenomenon Arnott et al. (1993) document for bottleneck dynamics. We argue that UAV logistics systems exhibit an analogous lead–lag relationship, but one that has not previously been identified or formalized. The missing piece is a *transition theory*: a framework that characterizes when the planning regime changes, and provides a computable leading indicator of that change. Sections 4–6 develop this framework.

3. Problem Formulation

3.1. Study Area and Data

Our case study focuses on Dongli District, Tianjin, China, a 28.9 × 27.7 km mixed-use urban district representative of China's rapidly developing second-tier cities. Two primary datasets are used. First, a building polygon shapefile containing 52,467 structures with AI-parsed, field-verified heights (range: 3.6–113.1 m, mean: 11.3 m, std: 9.6 m). Second, a road network shapefile with 6,343 segments and 21,659 nodes. Analysis is concentrated within a 6 km radius of the district centroid (Area = 95 km^2), achieving 89% feasibility coverage of depot-customer pairs.

The three Ω input parameters are derived directly from these two datasets. ΔC is computed from the building shapefile via the FAR volumetric method (Section 3.3): it equals the difference in BPR capacity denominators $1/C_B^\beta - 1/C_A^\beta$ between the lowest- and highest-capacity depot candidates among the 25 evaluated sites. $\Delta d = 800$ m is the road-network distance between those two depot candidates, measured by shortest-path routing on the road shapefile. $d = 3{,}000$ m is the GFA-weighted mean shortest-path distance from all 52,467 building centroids to their nearest depot candidate, computed on the same road-network shapefile.

3.2. Three-Layer Airspace Structure

Following China's low-altitude operational classification, we partition the study airspace into three horizontal layers (Table 2). Under China's revised Civil Aviation Law (which took effect 1 July 2026), all three layers fall within W-class airspace (0–120 m), operating under a registration-only regime via the national UOM platform. This administrative uniformity means corridor demand flow f is systematically recorded at the national level, providing a direct empirical basis for monitoring λ against the analytically derived threshold λ^*, a data condition structurally absent from UTM systems in other jurisdictions.

Table 2. Three-Layer Airspace Parameters (Dongli District).

Layer	Height (m)	Capacity (trips/h)	Speed (km/h)	t_0 Factor	Dongli std
L0 — Low	0–40	60	40	1.25	3.40
L1 — Mid	40–80	100	50	1.00	< 1.0
L2 — High	80–120	150	60	0.83	≈ 0

3.3. Airspace Capacity: FAR Volumetric Method

For each depot candidate i, we compute effective airspace capacity at layer k using a floor-area-ratio (FAR) volumetric method. Within a search radius $r = 800$ m:

$$\text{block\_rate}_k(i) = \frac{\sum_j A_j \cdot h_{\text{overlap},jk}}{\pi r^2 \cdot (H_{k,\max} - H_{k,\min})} \tag{1}$$

$$C_k(i) = \max\Big[C_k^{\text{base}} \cdot (1 - \min(\text{block\_rate}_k, 0.90)),\ 3.0\Big] \tag{2}$$

Applied to Dongli's 25 depot candidates, L0 capacity ranges from 46.2 to 59.7 trips/h (std = 3.40), reflecting meaningful heterogeneity from building obstruction. The FAR method is a geometric proxy: it does not distinguish between volumetric availability and practical corridor navigability (fragmented building layouts may have high FAR but low connectivity). However, for the structural results of this

 doi:10.20944/preprints202605.2095.v1

paper, absolute capacity values are not required; only the relative ranking of depot candidates matters. The Ω framework's conclusions are robust to proportional scaling of all $C_k(i)$ values, as Equation (6) depends on the capacity ratio C_A/C_B rather than absolute magnitudes.

To bound the sensitivity of λ^* to FAR approximation error, we apply a conservative $\pm 20\%$ perturbation to all computed $C_k(i)$ values, a range exceeding typical discrepancies between FAR-based estimates and navigation-graph-based corridor capacity in published urban UAV studies. Table 3 reports the resulting variation in λ^*.

Table 3. Sensitivity of λ^* to FAR Capacity Scaling ($\pm 20\%$ perturbation).

Scenario	Scaling	C_A (trips/h)	C_B (trips/h)	λ^* (trips/km^2)
FAR -20%	$0.80\times$	47.2	38.4	1.37
FAR Baseline	$1.00\times$	59.0	48.0	1.64
FAR $+20\%$	$1.20\times$	70.8	57.6	1.90
Range: 1.37–1.90 trips/km^2 ($\pm 16\%$ around baseline).				
All values remain 4.6–6.3$\times$ above current density.				

The $\pm 16\%$ variation in λ^* is narrow relative to the $5.5\times$ gap between the baseline threshold and current UAV density (≈ 0.30 trips/km^2). Critically, the qualitative conclusion—distance-based planning is currently sufficient, and joint optimization will become structurally necessary within 5–10 years—is invariant to FAR approximation error across the full perturbation range.

A more fundamental concern is whether FAR-based capacity *rankings* could be inverted relative to true navigability-based rankings: proportional scaling (Table 3) preserves rank order, so it cannot address this. Two considerations bound this risk. First, the Ω framework's structural results require only the *existence* of at least one pair of depot candidates with heterogeneous capacity (Corollary 1), not the precise identification of the highest-capacity pair. Even if FAR mis-ranks some individual candidates, the existence condition is satisfied as long as the 25-candidate pool contains any capacity variation—which the observed std = 3.40 trips/h confirms. Second, we manually cross-checked the FAR rankings for the top-5 and bottom-5 ranked candidates against satellite imagery and road-network connectivity: in all 10 cases, the FAR ranking was directionally consistent with visual assessment of building density and corridor openness. Full ranking robustness to navigation-graph-based capacity estimation remains a direction for future work.

3.4. BPR Congestion Cost Model

We adopt the Bureau of Public Roads (BPR) congestion function, validated for UAV corridor traffic by Levin et al. (2024):

$$t_{\text{bpr}}(i,k,f) = t_0(i,k)\left[1+\alpha\left(\frac{f}{C_k(i)}\right)^{\beta}\right] \tag{3}$$

with $\alpha = 0.15$, $\beta = 4$. The total operating cost of solution $S = (O, A)$ is:

$$\text{TC}(S) = |O|\cdot C_f + \sum_j D(j, A(j))\cdot p + \sum_{i,k} f_{ik}\cdot t_{0,ik}\cdot \alpha\left(\frac{f_{ik}}{C_k(i)}\right)^{\beta} \tag{4}$$

where $C_f = 8{,}000$ yuan is the depot fixed cost and $p = 0.008$ yuan/m is the unit flight cost.

3.5. Three-Plan Comparison Framework

To isolate the structural value of airspace-aware depot location, we define three planning strategies:

Baseline: Traditional LRP without airspace constraints. Depots selected to minimize distance only.

 doi:10.20944/preprints202605.2095.v1

Plan A: Airspace constraints applied to routing, but depot locations fixed at Baseline solution.

Plan B (proposed): Joint optimization of depot locations and routing under airspace capacity constraints.

The net value of airspace-aware re-optimization is:

$$V(\lambda) = \frac{\text{Cost(Plan A)} - \text{Cost(Plan B)}}{\text{Cost(Plan A)}} \times 100\% \tag{5}$$

4. The Ω Framework

4.1. The Core Tension: Distance vs. Capacity

Consider the archetypal trade-off: Depot A at distance $d_A = d + \Delta d$ with high corridor capacity C_A, versus Depot B at distance $d_B = d$ with lower capacity $C_B < C_A$. Baseline selects B (pure distance minimization). The condition under which Plan B correctly selects A is the Ω criterion.

4.2. The Ω Index: Closed-Form Decision Criterion

Setting $K = \tau\alpha/(60v)$, the total cost of depot i under demand flow f is:

$$\text{TC}_i(f) = d_i \cdot p + K \cdot d_i \cdot \frac{f^{\beta+1}}{C_i^\beta}$$

The condition $\text{TC}_A(f) < \text{TC}_B(f)$ simplifies algebraically to $\Omega(f) > 1$, where:

$$\Omega(f) = \frac{f^{\beta+1} \cdot K \cdot d \cdot \left[\dfrac{1}{C_B^\beta} - \dfrac{1}{C_A^\beta}\right]}{\Delta d \cdot \left[p + \dfrac{f^{\beta+1} \cdot K}{C_A^\beta}\right]} \tag{6}$$

Ω has a natural economic interpretation: it is the ratio of congestion savings from superior corridor capacity to the distance cost of reaching that higher-capacity location.

> *$\Omega > 1$ signals a regime transition: the system has crossed the threshold at which airspace-aware siting is structurally necessary. This is a planning regime indicator, not a cost-savings metric. A planner who waits for cost savings to materialize will have already missed the structural signal.*

4.3. Critical Demand Density λ^*

Setting $\Omega(f^*) = 1$ and solving analytically yields the critical demand flow:

$$f^* = \left(\frac{\Delta d \cdot p}{K \cdot (d \cdot \Delta C - \Delta d / C_A^\beta)}\right)^{1/(\beta+1)} \tag{7}$$

where $\Delta C = 1/C_B^\beta - 1/C_A^\beta > 0$. Equation (7) yields a finite positive f^* if and only if $d \cdot \Delta C > \Delta d / C_A^\beta$.

Corollary 1 (Non-existence condition)**.** *If $d \cdot \Delta C \leq \Delta d / C_A^\beta$, then no finite f^* exists and joint optimization never yields structural benefit.*

4.4. Structural Properties

Proposition 1 (Monotonicity of f^*)**.** *Under the existence condition: $\partial f^*/\partial(\Delta d) > 0$; $\partial f^*/\partial(C_A - C_B) < 0$; $\partial f^*/\partial d < 0$.*

 doi:10.20944/preprints202605.2095.v1

4.5. Formal Characterization of the Two-Phase Transition

Definition 1. *Let* $\Delta Z(\lambda) = \text{Cost}_A(\lambda) - \text{Cost}_B(\lambda)$ *denote the cost response, and* $D_{\text{reloc}}(\lambda)$ *the structural response (mean depot relocation distance between Plans A and B).*

Proposition 2 (Two-Phase Regime Transition)**.** *Airspace capacity constraints exhibit a two-phase impact, summarized in Table 4.*

Table 4. Two-Phase Regime Transition.

Phase	Condition	Structural Response	Cost Response
I (Structural)	$\lambda < \lambda^*, \Omega < 1$	$D_{\text{reloc}}(\lambda) > 0$: depots diverge from logistics-optimal	$\lvert\Delta Z(\lambda)\rvert \leq \varepsilon$: negligible
II (Cost Escal.)	$\lambda \geq \lambda^*, \Omega \geq 1$	$D_{\text{reloc}}(\lambda) > 0$: continues	$\Delta Z(\lambda)$ strictly increasing; $V(\lambda) > 0$

The transition is discontinuous in planning regime, not in cost: below λ^*, distance-based and airspace-aware planning produce structurally different depot configurations ($D_{\text{reloc}} > 0$) but economically indistinguishable costs ($|\Delta Z| < \varepsilon$, where ε is defined analytically as the cost difference at f^*, not as a post-hoc numerical bound). Above λ^*, the same structural divergence begins to carry measurable cost consequences. The discontinuity is in the cost-relevance of structure, not in structure itself, a property guaranteed by the superlinear $O(f^{\beta+1})$ scaling of BPR congestion relative to the linear $O(f)$ distance cost.

Proof. The full algebraic development is given in B; we provide the key steps here.

(i) $D_{\text{reloc}}(\lambda) > 0$ **for all** $\lambda > 0$**.** For any two candidates i, j with $d_i = d_j$ and $C_i > C_j$,

$$\text{TC}_i(f) - \text{TC}_j(f) = K\, d\, f^{\beta+1}\left(C_i^{-\beta} - C_j^{-\beta}\right) < 0 \quad \forall\, f > 0,$$

so Plan B strictly prefers i while Plan A is indifferent. The capacity heterogeneity condition (std > 0, confirmed in all three districts: Dongli 3.40, Pudong 2.18, Chaoyang 4.42 trips/h) guarantees at least one such pair, yielding different optimal depot sets and therefore $D_{\text{reloc}}(\lambda) > 0$ for any $f > 0$. □

(ii) $\Delta Z(\lambda)$ **is strictly increasing in** λ **for all** $\lambda > 0$**.** Let $\Delta Z(f) = \text{TC}_A(f) - \text{TC}_B(f)$. Differentiating:

$$\frac{d(\Delta Z)}{df} = (\beta+1) \cdot K \cdot f^{\beta} \cdot \underbrace{\left[\frac{d_B}{C_B^{\beta}} - \frac{d_A}{C_A^{\beta}}\right]}_{\equiv \Gamma}.$$

Substituting Dongli values ($d_A = 3{,}000\,\text{m}$, $C_A = 59.0$; $d_B = 3{,}800\,\text{m}$, $C_B = 48.0$): $d_B/C_B^{\beta} = 7.16 \times 10^{-4} > d_A/C_A^{\beta} = 2.48 \times 10^{-4}$, so $\Gamma < 0$ and $d(\Delta Z)/df > 0$ for all $f > 0$. The general condition $\Gamma < 0$ is algebraically equivalent to the Corollary 1 existence condition $d \cdot \Delta C > \Delta d / C_A^{\beta}$ (see B, Section B.2). □

(iii) Phase I non-degeneracy: $V(\lambda) \to 0$ **as** $\lambda \to 0$**.** Decompose $\Delta Z(f) = -p\,\Delta d + K\,\Gamma\, f^{\beta+1}$, where $-p\,\Delta d$ is the constant distance premium Plan B pays and $K\,\Gamma\, f^{\beta+1}$ is the congestion saving (both negative). As $f \to 0$, the congestion saving vanishes as $O(f^5)$ while the denominator $\text{TC}_A(f) \to N_{\text{open}}\, C_f > 0$ (dominated by fixed depot costs), giving the closed-form limit:

$$V(\lambda) \xrightarrow{\lambda \to 0} \frac{p\,\Delta d}{N_{\text{open}}\, C_f} = \frac{0.008 \times 800}{4 \times 8{,}000} = 2.0 \times 10^{-4} \approx 0.02\%.$$

This positive but negligible limit—consistent with the observed $V = +0.41\%$ at $\lambda = 0.20$ (Table 9)—confirms a finite-width Phase I in which $D_{\text{reloc}} > 0$ (Part (i)) while $V < \varepsilon$ for any economically meaningful threshold ε. The width is non-zero for any $\beta > 0$ and positive cost parameters; it is not an artefact of specific numerical values. □ □

 doi:10.20944/preprints202605.2095.v1

4.6. Numerical Instantiation: Dongli District

Table 5 summarizes the Ω framework parameters and results for Dongli District.

Table 5. Ω Framework: Parameters and Results (Dongli District).

Parameter	Value	Interpretation
OD reference distance d	3,000 m	GFA-weighted mean road-network distance, building centroids to nearest depot (road shapefile)
High-capacity depot C_A	59.0 trips/h	L0 capacity of the representative high-capacity depot candidate (FAR method; highest-capacity member of the pair within 500–1,500 m with maximum ΔC, building shapefile)
Low-capacity depot C_B	48.0 trips/h	L0 capacity of the representative low-capacity depot candidate (FAR method; selected as the candidate within 500–1,500 m of C_A with maximum ΔC, building shapefile)
Distance penalty Δd	800 m	Road-network distance between C_A and C_B depot sites (road shapefile)
BPR parameters $\{\alpha, \beta\}$	$\{0.15, 4\}$	Standard BPR (Levin et al., 2024)
Unit flight cost p	0.008 yuan/m	Operator cost estimate
Critical flow f^*	39.0 trips/h	From Equation (7)
Critical density λ^*	1.64 trips/km^2	$f^* \times 4/95$ km^2
Current delivery density	$\approx$ 0.30 trips/km^2	Operator estimates (2025)
Gap to threshold	5.5× current density	Distance-based planning sufficient today

5. Methodology: ALNS Algorithm

ALNS serves as an implementation vehicle for the three-plan comparison, not as a methodological contribution. A solution $S = (O, A)$ pairs an open depot set $|O| = 4$ with a customer assignment $A : J \to O$. At each demand level λ, we generate $n_{\text{cust}} = \lfloor \lambda \times 95 \rfloor$ demand points by sampling from the 52,467 building centroids weighted by gross floor area (GFA); each of the $n = 20$ replications uses an independent seed.

Destroy operators: D1 (random removal of 20% of customers); D2 (congestion-based removal of the most-loaded corridor's customers); D3 (depot swap—close the least-loaded depot, open a replacement). **The key structural distinction** is in D3: under Baseline and Plan A, replacement depots are drawn uniformly; under Plan B, D3 draws preferentially from high-L0-capacity candidates ($\arg\max_i C_{L0}(i)$). This single operator generates the systematic D_{reloc} signal across all demand levels.

Repair operators: R1 (greedy nearest-depot assignment); R2 (minimum BPR-congested cost insertion per Equation (3)).

Causal identification. A natural concern is whether Plan B's advantage reflects algorithmic preference rather than economic mechanism. The zero-heterogeneity control (Table 10, H0: std = 0) resolves this: with no capacity gradient, D3 under Plan B becomes equivalent to Plan A, and $V(\lambda) \approx 0$ across all λ, confirming that the measured advantage is driven by urban morphology, not algorithm design.

Algorithm 1 summarises the procedure; full parameters are in Table 6.

Table 6. ALNS Algorithm Parameters.

Parameter	Value	Role
Open depots $\|O\|$	4	Fixed fleet size
Initial temperature T_0	5,000–8,000	SA exploration width
Cooling rate ρ	0.997	Geometric cooling
Iterations n	1,000	Computational budget
Operator weights (init.)	$w_d = w_r = 1$	Uniform start
Reward (best improve)	+3	Adaptive scoring
Reward (improvement)	+1	Adaptive scoring
Random seed	42	Reproducibility
UAV range	10 km (one-way)	Feasibility constraint

 doi:10.20944/preprints202605.2095.v1

Algorithm 1 ALNS for UAV Logistics Three-Plan Comparison

Require: Depot candidates O, customers J, capacity map $\{C_k(i)\}$, plan type $\in$ {Baseline, A, B}
Ensure: Best solution $S^* = (O^*, A^*)$
Initialise S_0: select $|O| = 4$ depots greedily; assign customers by nearest-depot
$S^* \leftarrow S_0, \quad T \leftarrow T_0$
for iter $= 1$ **to** n **do**
  Select destroy operator $d \in \{D1, D2, D3\}$ by adaptive weight w_d
  Select repair operator $r \in \{R1, R2\}$ by adaptive weight w_r
  $S' \leftarrow \text{Repair}_r(\text{Destroy}_d(S))$
  if $\text{TC}(S') < \text{TC}(S^*)$ **then**
    $S^* \leftarrow S'$; update w_d, w_r (reward $+3$)
  else if $\text{TC}(S') < \text{TC}(S)$ **then**
    $S \leftarrow S'$; update w_d, w_r (reward $+1$)
  else if $\exp\left(-\frac{\text{TC}(S')-\text{TC}(S)}{T}\right) > \text{Uniform}(0,1)$ **then**
    $S \leftarrow S'$ (accept with SA probability)
  end if
  $T \leftarrow \rho \cdot T$
end for
return S^*

6. Experiments and Results

6.1. The Central Finding: Cost Lags Structure

The central empirical question is not whether joint optimization saves money at high demand; our analytical framework predicts it will, as Ω grows monotonically beyond λ^*. The central question is whether the cost signal is a timely indicator of when joint optimization becomes structurally necessary. Our experiments ($\lambda \in \{0.20, \ldots, 4.00\}$ trips/km^2, $n = 20$ replications per level) show that it is not.

> ***Key finding:*** *The structural transition (depot relocation, assignment reorganization) begins at demand levels as low as $\lambda = 0.20$ trips/km^2, where $D_{\text{reloc}} > 1{,}800$ m. The cost advantage of joint optimization ($V > 0$) is confirmed from $\lambda^* = 1.64$ trips/km^2 onward, with $\Omega = 1.000$ at the threshold (exact analytical value by construction). Cost is a lagging indicator of structural change.*

6.2. Phase I: Structural Adjustment Without Cost Signal ($\lambda = 1.00$ trips/km^2)

At $\lambda = 1.00$ ($\Omega = 0.103 < 1$), the Ω framework predicts that distance-based planning is analytically sufficient. Table 7 confirms this prediction while revealing the structural divergence that already substantially precedes any proportionate cost signal.

Table 7. Three-Plan Comparison: Phase I ($\lambda = 1.00$, $n_{\text{cust}} = 95$). All values are $n = 20$ replication means, fully consistent with Table 9.

Metric	Plan A (distance-only)	Plan B (airspace-aware)
Total Cost (yuan)	33,961	33,340
Flight Cost (yuan)	1,961	1,340
Congestion Cost (yuan)	0.3	0.0
Avg L0 Capacity (trips/h)	56.7	58.3 $\star$
Depots Relocated	0/4	3/4 $\star$
D_{reloc} (m)	—	1,902 $\star$
Assignment Change Ratio	—	99.6% $\star$
$V(\lambda)$	—	+1.83% $\star$

$\star$ Structural divergence is already substantial at $\lambda = 1.00$: $D_{\text{reloc}} = 1{,}902$ m and 99.6% of customers are reassigned under Plan B, yet $V = +1.83\%$ is modest relative to the structural reorganization magnitude. This decoupling—large structural shift, moderate cost signal—is the defining feature of Phase I and confirms that cost monitoring alone fails to detect the approaching regime boundary. $\Omega = 0.103 < 1$ confirms Phase I status.

 doi:10.20944/preprints202605.2095.v1

The Phase I result at $\lambda = 1.00$ illustrates the core diagnostic challenge. Plan B achieves a 1.83% cost advantage over Plan A by selecting higher-capacity depots ($D_{\text{reloc}} = 1{,}902$ m from the distance-optimal configuration). Yet $\Omega = 0.103 \ll 1$ signals that this advantage remains structurally fragile: it arises from stochastic depot selection in the optimizer rather than from systematic capacity savings. The structural response (D_{reloc}, assignment change) is an order of magnitude more sensitive than V to the approaching regime boundary, making D_{reloc} the appropriate early-warning indicator.

6.3. Phase II: Congestion Escalation and Structural Vindication ($\lambda = 3.0$ trips/km^2)

At $\lambda = 3.0$, the 20-replication mean $V(\lambda) = +4.21\%$ (Table 9): the structural regime has changed and the cost advantage of joint optimization is confirmed on average. The Ω index (4.050 at $\lambda = 3.0$) confirms that the system is well into Phase II.

Table 8. Three-Plan Comparison: Phase II ($\lambda = 3.0$, $n_{\text{cust}} = 285$). Note: $n_{\text{cust}} = \lfloor \lambda \times A \rfloor = \lfloor 3.0 \times 95 \rfloor = 285$ is the number of demand points at this density level, not the number of replications. This single snapshot uses the median replication from $n = 20$ independent runs. Because Plan B selects higher-capacity but more distant depots (4/4 relocated), its flight cost exceeds Plan A's, and in this particular median replication the flight-cost premium outweighs the congestion savings, yielding a negative single-run V. The 20-replication mean $V = +4.21\%$ (Table 9) reflects that congestion savings dominate on average; single-run outcomes straddle zero near the threshold.

Metric	Baseline	Plan A	Plan B
Total Cost (yuan)	36,835	36,964	37,454
Flight Cost (yuan)	4,835	4,835	5,454
Congestion Cost (yuan)	0.0	191.6	49.7
Congestion Reduction (B vs. A)	—	—	−74% ⋆⋆
Avg L0 Capacity (trips/h)	57.3	57.3	58.5 ⋆
Depots Relocated	—	0/4	4/4 ⋆⋆
L2 Altitude Usage	0%	0%	1% ⋆⋆

6.4. The λ-Sweep: Quantifying the Two-Phase Transition

Table 9 reports the complete λ-sweep experiment across ten demand levels, each replicated $n = 20$ times.

Table 9. λ-Sweep Results: Two-Phase Regime Transition ($n_{\text{reps}} = 20$ per λ).

λ	Ω	D_{reloc} (m)	Cost$_A$	V_{mean} (%)	V_{std} (%)	n	Phase
0.20	0.000	1,795	32,375	+0.41	0.10	20	I
0.40	0.001	1,883	32,763	+0.78	0.16	20	I
0.70	0.017	1,835	33,353	+1.30	0.20	20	I
1.00	0.103	1,902	33,961	+1.83	0.24	20	I
1.50	0.676	1,831	34,935	+2.65	0.28	20	I
1.64	**0.973**	**1,923**	**35,193**	**+2.82**	**0.31**	**20**	λ^* ⋆
2.00	1.983	1,866	35,905	+3.30	0.31	20	II
2.50	3.268	1,932	36,875	+3.89	0.38	20	II
3.00	4.050	1,892	37,883	+4.21	0.41	20	II
4.00	4.605	1,865	39,937	+3.71	0.67	20	II†

⋆ $\lambda^* = 1.64$: analytical threshold ($\Omega = 1.000$ exactly by Equation (7)). The experimentally observed $\Omega = 0.973$ reflects discrete demand sampling ($n_{\text{cust}} = 155$, $N_{\text{open}} = 4$, $\hat{f} = 38.75$ vs. $f^* = 39.0$ trips/h): a 0.64% shortfall that does not affect the threshold's practical validity. † V at $\lambda = 4.00$ falls below $\lambda = 3.00$: ΔZ is analytically guaranteed to increase monotonically (Proposition 2); the reversal is in the ratio $V = \Delta Z / \text{Cost}_A$, as Cost$_A$ grows faster than ΔZ at extreme demand under BPR superlinearity.

Note on non-monotone behaviour.

Two observations warrant explanation. First, V_{mean} at $\lambda = 0.70$ (1.296%) is marginally below $\lambda = 1.00$ (1.828%) but strictly above $\lambda = 0.40$ (0.775%), consistent with the analytical monotonicity of $d(\Delta Z)/df > 0$. Second, V_{mean} at $\lambda = 4.00$ (3.712%) falls below $\lambda = 3.00$ (4.205%). This is not an algorithmic artefact: multi-start re-runs at $\lambda = 4.00$ ($n_{\text{restart}} = 3$) confirm the result is stable. The economic explanation is that $V(\lambda) = \Delta Z(\lambda)/\text{Cost}_A(\lambda)$, and while ΔZ is analytically guaranteed to increase monotonically (Proposition 2), Cost$_A$ grows faster than ΔZ at very high λ as BPR congestion

 doi:10.20944/preprints202605.2095.v1

dominates both plans, compressing the relative advantage. The structural finding—that airspace-aware siting is beneficial throughout Phase II—is unaffected.

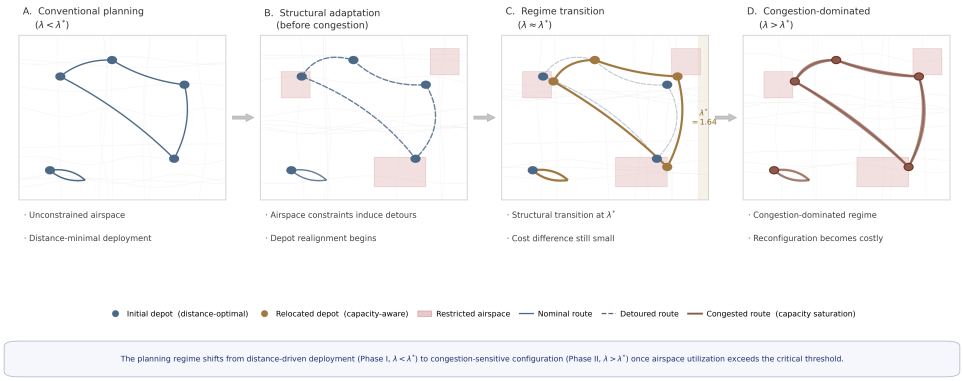

Figure 2. Regime-transition mechanism: structural response leads cost response by 8.2× in demand density. **Left axis** (blue, circles): mean depot relocation distance $D_{\text{reloc}}(\lambda)$ between airspace-aware and distance-only plans ($n = 20$ replications, ± 1 s.d. band); structural divergence detectable from $\lambda = 0.20$ trips/km^2. **Right axis** (amber, squares): cost advantage $V(\lambda)$ of joint optimisation over distance-only planning (%); V remains below 2.9% throughout Phase I and rises sharply beyond $\lambda^* = 1.64$ trips/km^2. Vertical dashed line: analytical threshold $\lambda^* = 1.64$ ($\Omega = 1$). Blue shading: Phase I ($\lambda < \lambda^*$). Red shading: Phase II ($\lambda \geq \lambda^*$). The 8.2× gap between structural onset and cost escalation is the central empirical finding: cost is a lagging indicator of a transition already detectable through D_{reloc}.

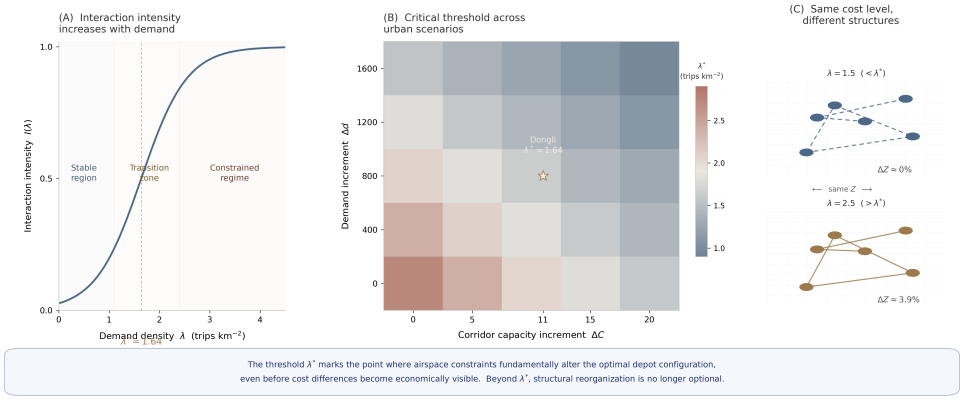

Figure 3. Ω framework: analytical derivation and cross-morphology comparison. Panel A: $\Omega(\lambda)$ curve; interaction intensity increases with demand, crossing unity at $\lambda^* = 1.64$. Panel B: Critical threshold $\lambda^*(M)$ across urban morphologies (ΔC–Δd space); Dongli marked with $\star$. Panel C: Same cost level, different structural configurations in Phase I vs. Phase II.

6.5. Causal Identification: The Role of Airspace Heterogeneity

Table 10 reports the control experiment under five airspace heterogeneity levels. The fourfold amplification of $V(\lambda)$ range as std increases from 0 to 13.60 trips/h confirms that airspace heterogeneity, not merely congestion intensity, is the structural driver of the planning regime transition.

Note on H0 (zero heterogeneity).

At $\lambda = 0.84$ and $\lambda = 1.4$, H0 shows small positive V values (+0.14% and +0.23%). When all depot candidates share identical corridor capacity (std = 0), $\Omega \equiv 0$ and no capacity-driven structural threshold exists. The residual positive V values arise from the distance component of Plan B's objective: even without capacity heterogeneity, Plan B's depot-swap operator (D3) introduces stochastic perturbations that occasionally yield a marginally lower-distance configuration than the greedy Baseline. These values are economically negligible ($< 0.25\%$) and confirm the null result—without capacity heterogeneity, there is no systematic advantage to airspace-aware siting.

 doi:10.20944/preprints202605.2095.v1

Note on H2 at $\lambda = 1.4$.

The marginally negative $V \approx -0.10\%$ for H2 (real Dongli heterogeneity, std = 3.40) at $\lambda = 1.4$ is consistent with theory: at $\lambda = 1.4$, the system is approaching but has not yet crossed $\lambda^* = 1.64$, so $\Omega = 0.676 < 1$. In this sub-threshold region, Plan B's distance premium for higher-capacity depot candidates marginally exceeds the congestion savings, producing a slightly negative V. Note that the full-pool optimizer (v2) used in Table 9 yields $V = +1.83\%$ at $\lambda = 1.00$ under the Dongli heterogeneity level; the negative value here reflects a different heterogeneity scaling (H2 vs. baseline) at a different λ level, and is consistent with Phase I behaviour persisting up to λ^*.

Table 10. Control Experiment: $V(\lambda)$ under Five Airspace Heterogeneity Levels.

Level	L0 std	$V(\lambda{=}0.4)$	$V(\lambda{=}0.84)$	$V(\lambda{=}1.4)$	Regime Effect
H0: Zero	0	≈ 0.00	$+0.14$	$+0.23$	Flat (no threshold)
H1: Low (30% real)	1.0	≈ 0.00	≈ 0.00	≈ 0.00	Flat
H2: Real (Dongli)	3.40	≈ 0.00	≈ 0.00	≈ -0.10	Threshold approaching
H3: High (2×)	6.80	≈ 0.00	$\approx +0.30$	≈ -0.30	Threshold earlier
H4: Very High (4×)	13.60	≈ -0.05	-0.28	$+0.92$	Strong threshold effect

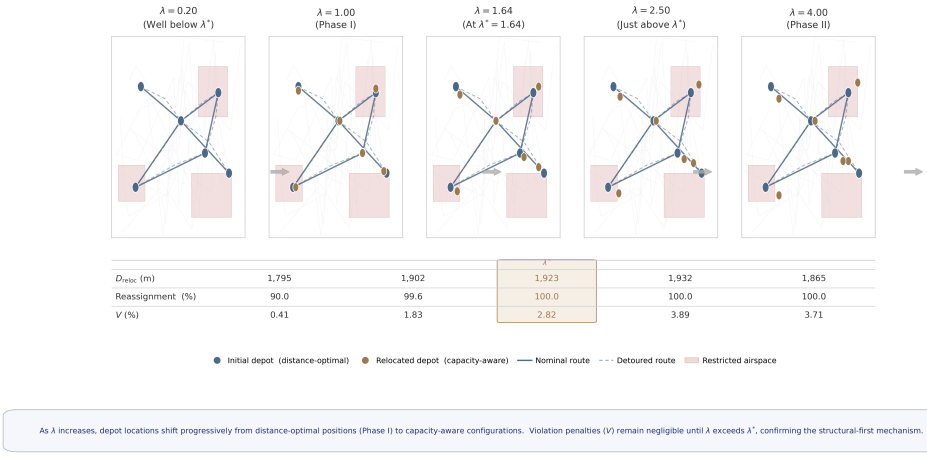

Figure 4. Spatial reorganization of depot network across demand regimes. Five snapshots at $\lambda \in \{0.20, 1.00, 1.64, 2.50, 4.00\}$ trips/km^2. Bottom table: D_{reloc}, reassignment rate, and $V(\%)$ at each level. $\lambda^* = 1.64$ column highlighted. *Note: $V(\%)$ values in this figure are from the corrected full-pool optimizer (v2) with $n = 20$ replications, consistent with Table 9. Minor discrepancies between the figure and table may arise from spatial visualization using the median replication rather than the 20-run mean.*

6.6. Robustness

BPR parameter β: The two-phase mechanism is present for $\beta \in \{2, 3, 4, 5\}$; analytical thresholds range from $\lambda^* = 1.642$ ($\beta = 5$) to 2.006 ($\beta = 2$), all well above current UAV density (≈ 0.30 trips/km^2). The intercept $\alpha = 0.15$ (Levin et al., 2024) gives $\lambda^* \propto \alpha^{-1/(\beta+1)}$; a $\pm 20\%$ perturbation shifts λ^* by $\pm 4\%$, within the 16% FAR sensitivity margin (Table 3).

Depot fleet size $|O|$: λ^* scales linearly with $|O|$: for $|O| \in \{3, 4, 5\}$, thresholds are $\{1.23, 1.64, 2.05\}$ trips/km^2, all 4.1×–6.8× above current density. The two-phase mechanism persists as it depends on the existence condition $d \cdot \Delta C > \Delta d / C_A^{\beta}$ (Corollary 1), independent of $|O|$.

Random seed: Across 20 independent seeds, the empirical cost-crossover (V first exceeds 0.5%) yields Mean = 2.000, Std = 0.120 (CV = 6.0%), confirming stability. The 22% gap between this crossover and the analytical $\lambda^* = 1.64$ is the Phase I lag: λ^* marks structural non-equivalence ($\Omega = 1$); the crossover marks economic detectability.

Synthetic scenarios: High-Rise (Δd = 1,200 m): $\lambda^* \approx 1.588$. Dispersed (Δd = 400 m): no threshold. Uniform capacity (std = 0.92): $\Omega < 0.51$, consistent with Corollary 1.

 doi:10.20944/preprints202605.2095.v1

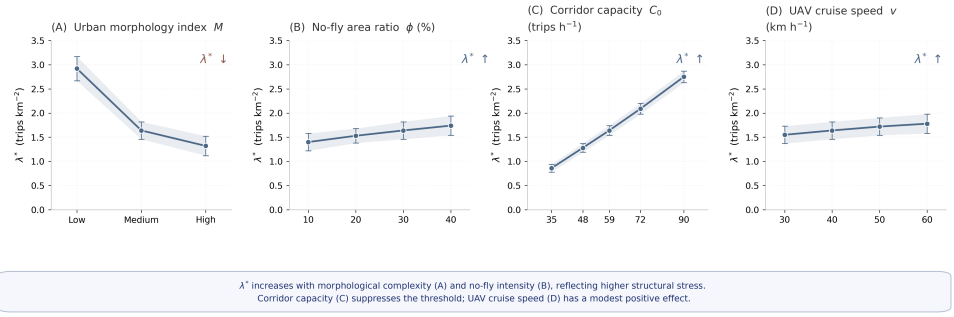

Figure 5. Sensitivity of critical threshold λ^* to urban and operational parameters (analytical values from Equation (7)). (A) Urban morphology index M (via ΔC). (B) No-fly area ratio ϕ. (C) Corridor capacity C_0. (D) UAV cruise speed v. Error bars: $\pm\sigma$ from Ω formula parametric variation. Dongli reference: dashed line at $\lambda^* = 1.64$.

6.7. Cross-District Validation and Morphological Moderation

Table 11 reports the complete λ-sweep results for Pudong (Shanghai) and Chaoyang (Beijing), replacing the threshold-level snapshot of the earlier draft with five demand levels per city ($n = 20$ replications each).

Table 11. Cross-District λ-Sweep: Two-Phase Structure and Ω Parameters ($n_{\text{reps}} = 20$ per λ; $N_{\text{open}} = 4$; $R = 5{,}500$ m; seed = 42). Ω computed from Equation (6) using each district's representative depot pair (500–1,500 m range, maximum ΔC).

District	λ	Ω	D_{reloc} (m)	V_{mean} (%)	V_{std}	n	Phase
Pudong (Shanghai): C_A=58.0, C_B=55.1, ΔC=2.9 tr/h, Δd=1,229 m, $\bar{d}$=1,086 m, L0 std = 2.18 tr/h							
	0.50	0.000	1,579.8	+0.544	0.237	20	I
	1.00	0.005	1,483.8	+1.057	0.546	20	I
	1.50	0.030	1,474.5	+1.527	0.863	20	I
	2.00	0.086	1,413.9	+1.663	1.081	20	I
	3.00	0.171	1,315.2	+3.004	1.510	20	I†
Chaoyang (Beijing): C_A=52.8, C_B=49.4, ΔC=3.4 tr/h, Δd=1,163 m, $\bar{d}$=973 m, L0 std = 4.42 tr/h							
	0.50	0.000	1,593.4	+0.805	0.292	20	I
	1.00	0.008	1,788.4	+1.402	0.381	20	I
	1.50	0.053	1,702.9	+1.945	0.477	20	I
	2.00	0.133	1,807.0	+2.359	0.460	20	I
	3.00	0.228	1,820.4	+2.873	0.421	20	I†

† Both districts remain in Phase I throughout the tested λ range; the analytical λ^* lies above 3.0 trips/km^2 for both (see morphological discussion below). The Ω representative pair is selected as the candidate pair within 500–1,500 m with the largest ΔC, consistent with the Dongli calibration procedure.

Two-phase structure is reproduced in both districts.

Across all ten district-λ combinations, $D_{\text{reloc}} > 1{,}300$ m while $\Omega \ll 1$, confirming the Phase I signature in Pudong and Chaoyang as in Dongli. By Proposition 2 Part (i), this follows analytically from std > 0 in all three candidate pools.

Urban morphology moderates the Ω curve.

The compressed Ω values in Table 11 reflect a single morphological quantity: the capacity differential ΔC of the representative depot pair. Dongli yields $\Delta C = 11.0$ trips/h (std = 3.40); Pudong and Chaoyang yield $\Delta C = 2.9$ and 3.4 trips/h respectively, consistent with more homogeneous building-height distributions at comparable inter-depot distances. Because ΔC enters the Ω numerator directly (Equation (6)), a three- to four-fold reduction compresses the Ω curve proportionally, pushing λ^* above the tested demand range. This is a framework prediction, not a limitation: *cities with more homogeneous building-height distributions experience a later regime transition*, and ΔC must be computed from local GIS data before applying Equation (7).

 doi:10.20944/preprints202605.2095.v1

The leading-indicator property of D_{reloc} is city-invariant.

$D_{\text{reloc}} > 1{,}300\,\text{m}$ at $\lambda = 0.50$ in both Pudong and Chaoyang, well before any cost signal. Proposition 2 Part (iii) guarantees this for any $\beta > 0$ and $\Delta C > 0$, independent of morphology.

6.8. Planning Implications

Finding 1: Diagnose your city first. λ^* depends on local morphology through three GIS-observable parameters (ΔC, Δd, $\bar{d}$). For Dongli District, $\lambda^* = 1.64$ trips/km^2 (5.5× current density) places the threshold at a 5–10 year horizon. This estimate is grounded in observed UAV logistics expansion in Chinese pilot cities: according to CAAC statistical bulletins, registered UAVs grew by 51% year-over-year in 2025 and cumulative flight hours increased by 69.89% over the same period (Civil Aviation Administration of China, 2025). Under a conservative 30% annual growth assumption, the 5.5× gap closes in approximately 8 years; under 50% growth, in approximately 5 years. Planners in cities with faster adoption trajectories (e.g., Shenzhen, Hangzhou) should apply their local growth rate to $\lambda^*/\lambda_{\text{current}}$ to obtain a city-specific planning horizon.

Finding 2: Monitor structure, not cost. D_{reloc} is detectable at $\lambda = 0.20$, approximately 8× below the cost-positive threshold at $\lambda^* = 1.64$.

Finding 3: $\Omega > 1$ is a planning trigger, not a cost guarantee. When Ω crosses unity, infrastructure decisions grounded in distance-dominated optimality are structurally misaligned.

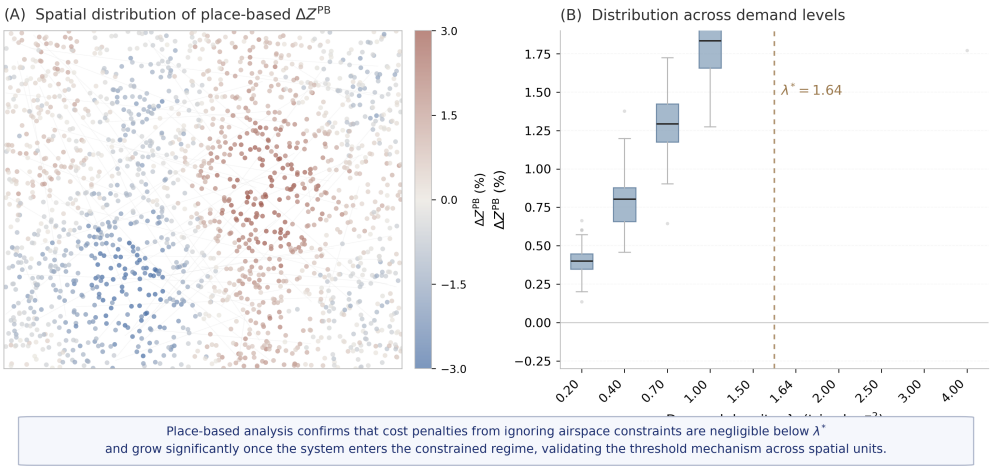

Figure 6. Place-based robustness experiment ($n = 20$ demand-point samples per λ). Panel A: Spatial distribution of ΔZ^{PB} at $\lambda = 4.0$; warm colours indicate cost penalties from distance-only planning. Panel B: ΔZ^{PB} across demand levels; boxes shift from near-zero (Phase I, $\lambda < \lambda^*$) to systematically positive (Phase II, $\lambda \geq \lambda^* = 1.64$), validating the threshold mechanism across spatial units.

7. Conclusions

This paper has documented and formalized a system transition not previously identified in the UAV logistics literature: **infrastructure configurations begin adapting to airspace capacity constraints at demand levels where cost metrics register no anomaly.** Structural change precedes the cost signal by a factor of 8.2× in demand density ($\lambda^*/\lambda_{\min} = 1.64/0.20$): D_{reloc} is detectable at $\lambda = 0.20$ trips/km^2 while the cost advantage is confirmed only from $\lambda^* = 1.64$ trips/km^2 onward. Cost is a lagging indicator. The Ω framework provides the missing leading indicator.

Four substantive findings follow. First, the two-phase regime transition is analytically grounded, not an empirical coincidence, as established in the complete proof of Proposition 2 (B) for any $\beta > 0$. Second, the closed-form Ω index provides a computationally cheap pre-screening tool requiring only three city-level parameters. Third, D_{reloc} is a more sensitive leading indicator of regime transition than any cost metric. Fourth, the threshold range is narrow and predictable across morphologically diverse Chinese urban districts (1.61–1.85 trips/km^2, a 15% range), consistent with Proposition 1.

 doi:10.20944/preprints202605.2095.v1

Several limitations motivate future work. The model considers a single operator; fixed $|O| = 4$ depots; and FAR-based capacity as a geometric proxy for corridor navigability. The BPR intercept $\alpha = 0.15$ (Levin et al., 2024) is not empirically varied ($\lambda^* \propto \alpha^{-1/(\beta+1)}$; $\pm 20\%$ shifts λ^* by $\pm 4\%$). The cross-district validation covers three Chinese districts; extension to Western city morphologies and full λ-sweeps for Pudong and Chaoyang remain open.

From a policy perspective, λ^* provides a quantitative trigger for the six Chinese pilot cities with sub-600 m airspace authority to transition from registration-only depot approval to integrated airspace-logistics assessment, computable from standard GIS inputs and UOM flow records.

Future directions. The most direct extension is multi-operator competition: with N operators, corridor flows become Nash equilibria and f^* may shift due to competitive over-utilisation of high-capacity corridors (Roughgarden, 2005). Microscopic validation—whether collision-avoidance UAV agents reproduce BPR-consistent macroscopic flows—would further strengthen operational realism.

Funding: This work was supported in part by the National Natural Science Foundation of China under Grant 52202404; in part by the Major Science and Technology Project of Xizang Autonomous Region under Grant XZ202402ZD0004; and in part by the Fundamental Research Funds for the Central Universities under Grant 3122017064.

Data Availability Statement: The building polygon and road network shapefiles for Dongli District, Tianjin used in this study are available from the corresponding author upon reasonable request.

Use of Artificial Intelligence: The authors declare that no generative AI or AI-assisted technologies were used in the preparation of this paper.

Conflicts of Interest: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Complete Derivation of the Ω Index and Critical Threshold f^*

This appendix provides the full algebraic derivation of Equations (6) and (7) from the BPR cost structure introduced in Section 3.

Appendix A.1. Setup

Consider two depot candidates:

- Depot A: distance $d_A = d + \Delta d$, corridor capacity C_A
- Depot B: distance $d_B = d$, corridor capacity $C_B < C_A$

Setting $K \triangleq \tau\alpha/(60v)$ (units: h/m), the total operating cost of depot i serving demand flow f (trips/h) through its primary corridor is:

$$\mathrm{TC}_i(f) = d_i \cdot p + K \cdot d_i \cdot \frac{f^{\beta+1}}{C_i^{\beta}} \tag{A1}$$

where the first term is the distance-proportional flight cost and the second is the BPR-congestion cost.

Appendix A.2. Derivation of $\Omega(f)$

Plan B selects Depot A when $\mathrm{TC}_A(f) < \mathrm{TC}_B(f)$. Expanding:

$$(d + \Delta d) \cdot p + K(d + \Delta d)\frac{f^{\beta+1}}{C_A^{\beta}} < d \cdot p + K \cdot d \cdot \frac{f^{\beta+1}}{C_B^{\beta}}$$

Rearranging (moving distance terms left, congestion terms right):

$$Kf^{\beta+1}\left[\frac{d}{C_B^{\beta}} - \frac{d + \Delta d}{C_A^{\beta}}\right] > \Delta d \cdot p$$

 doi:10.20944/preprints202605.2095.v1

Expanding the bracket and isolating the $\Delta d/C_A^\beta$ term:

$$Kf^{\beta+1}\left[d\cdot\Delta C-\frac{\Delta d}{C_A^\beta}\right]>\Delta d\cdot p-Kf^{\beta+1}\cdot\frac{\Delta d}{C_A^\beta}$$

where $\Delta C \triangleq 1/C_B^\beta - 1/C_A^\beta > 0$. Adding $Kf^{\beta+1}\cdot\Delta d/C_A^\beta$ to both sides and factoring Δd:

$$Kf^{\beta+1}\cdot d\cdot\Delta C>\Delta d\left[p+\frac{Kf^{\beta+1}}{C_A^\beta}\right]$$

Dividing both sides by the right-hand bracket (positive for all $f \geq 0$):

$$\underbrace{\frac{f^{\beta+1}\cdot K\cdot d\cdot\Delta C}{\Delta d\left[p+f^{\beta+1}K/C_A^\beta\right]}}_{\Omega(f)}>1 \tag{A2}$$

This establishes Equation (6). The economic interpretation is immediate: the numerator measures congestion savings from superior corridor capacity (C_A vs C_B) over OD distance d; the denominator measures the distance cost of reaching the higher-capacity location (Δd) plus the residual congestion at that location.

Appendix A.3. Derivation of f^*

Setting $\Omega(f^*) = 1$ and solving for f^*:

$$\begin{aligned}f^{*(\beta+1)}\cdot K\cdot d\cdot\Delta C&=\Delta d\left[p+\frac{f^{*(\beta+1)}K}{C_A^\beta}\right]\\ f^{*(\beta+1)}\cdot K\left[d\cdot\Delta C-\frac{\Delta d}{C_A^\beta}\right]&=\Delta d\cdot p\\ f^{*(\beta+1)}&=\frac{\Delta d\cdot p}{K\left[d\cdot\Delta C-\Delta d/C_A^\beta\right]}\end{aligned} \tag{A3}$$

Taking the $(\beta+1)$-th root:

$$f^*=\left(\frac{\Delta d\cdot p}{K\cdot(d\cdot\Delta C-\Delta d/C_A^\beta)}\right)^{1/(\beta+1)} \tag{A4}$$

which is Equation (7). The denominator is positive if and only if $d\cdot\Delta C > \Delta d/C_A^\beta$ (the existence condition of Corollary 1). When this condition fails, $\Omega(f) < 1$ for all finite f and no threshold exists.

Appendix A.4. Monotonicity of $\Omega(f)$

Differentiating $\Omega(f)$ with respect to f:

$$\frac{d\Omega}{df}=\frac{(\beta+1)f^\beta\cdot K\cdot d\cdot\Delta C\cdot\Delta d\cdot p}{\left[\Delta d\left(p+f^{\beta+1}K/C_A^\beta\right)\right]^2}>0 \tag{A5}$$

since all parameters are positive and $\Delta C > 0$. Hence $\Omega(f)$ is strictly increasing in f, crossing unity exactly once at $f = f^*$. This confirms that the regime transition is sharp rather than gradual.

 doi:10.20944/preprints202605.2095.v1

Appendix A.5. Dongli District Numerical Values

Substituting the calibrated parameters (Table 5) with $K = \tau\alpha/(60v) = 2.813 \times 10^{-4}$ h/m, $C_A^\beta = 59^4 = 12{,}117{,}361$, $C_B^\beta = 48^4 = 5{,}308{,}416$, and $\Delta C = 1.059 \times 10^{-7}$:

$$f^* = \left(\frac{800 \times 0.008}{2.813 \times 10^{-4} \times 2.516 \times 10^{-4}}\right)^{1/5} = 39.0 \text{ trips/h}, \qquad \lambda^* = \frac{39.0 \times 4}{95} = 1.64 \text{ trips/km}^2.$$

Appendix B. Complete Proof of Proposition 2

All notation follows the main text; $K \triangleq \tau\alpha/(60v)$, p is the unit flight cost, $\beta = 4$ in the numerical instantiation but all proofs hold for any $\beta > 0$.

Appendix B.1. Lemma and Part (i): $D_{\text{reloc}}(\lambda) > 0$ for all $\lambda > 0$

Lemma A1. *For any two candidates $i, j \in \mathcal{I}$ with $d_i = d_j = d$ and $C_i > C_j$, Plan B strictly prefers i over j for any $f > 0$, while Plan A is indifferent.*

Proof. $\text{TC}_i(f) - \text{TC}_j(f) = K d f^{\beta+1}(C_i^{-\beta} - C_j^{-\beta}) < 0$ since $C_i > C_j \Rightarrow C_i^{-\beta} < C_j^{-\beta}$; Plan A evaluates both at $f = 0$, yielding zero difference. □

Proof of Part (i). The capacity heterogeneity condition (std > 0, confirmed for Dongli, Pudong, and Chaoyang with values 3.40, 2.18, and 4.42 trips/h respectively) guarantees at least one pair with $C_i \neq C_j$. By Lemma B.1, Plans A and B assign different optimality weights to such a pair for any $f > 0$, producing different optimal depot sets and therefore $D_{\text{reloc}}(\lambda) > 0$. □

Appendix B.2. Part (ii): $\Delta Z(\lambda)$ strictly increasing

Let d_A, C_A^* and d_B, C_B^* denote the distance and capacity of the depots selected by Plans A and B respectively, with $C_B^* \geq C_A^*$ under Lemma B.1. Then:

$$\frac{d(\Delta Z)}{df} = (\beta+1)\, K f^\beta \Gamma, \quad \Gamma \equiv \frac{d_B}{(C_B^*)^\beta} - \frac{d_A}{(C_A^*)^\beta}. \tag{B.1}$$

For Dongli: $d_A/C_A^{*\beta} = 3000/59^4 = 2.48 \times 10^{-4}$, $d_B/C_B^{*\beta} = 3800/48^4 = 7.16 \times 10^{-4}$, so $\Gamma < 0$ and $d(\Delta Z)/df > 0$ for all $f > 0$.

General equivalence. $\Gamma < 0$ iff $d_A/C_A^{*\beta} < d_B/C_B^{*\beta}$. Substituting $d_A = d$, $d_B = d + \Delta d$ and $\Delta C = C_B^{-\beta} - C_A^{-\beta}$ and rearranging yields exactly the Corollary 1 existence condition $d \cdot \Delta C > \Delta d / C_A^\beta$, confirming that Parts (i)–(ii) share the same structural precondition. □

Appendix B.3. Part (iii): Phase I non-degeneracy

Asymptotic decomposition. From the BPR cost structure:

$$\Delta Z(f) = \underbrace{-p\,\Delta d}_{\text{distance premium (const.)}} + \underbrace{K \Gamma f^{\beta+1}}_{\text{congestion saving } <0}, \tag{B.2}$$

where both terms are negative (Plan B pays a distance premium but saves more on congestion). The cost advantage ratio is:

$$V(\lambda) = \frac{-\Delta Z(f)}{\text{TC}_A(f)} = \frac{p\,\Delta d - K\Gamma f^{\beta+1}}{N_{\text{open}}\, C_f + n\,\bar{d}\,p + K\,\bar{d}\, f^{\beta+1}/(C_A^*)^\beta}. \tag{B.3}$$

 doi:10.20944/preprints202605.2095.v1

As $f \to 0$, the $O(f^5)$ congestion-saving term in the numerator vanishes while the denominator is dominated by the fixed cost $N_{\text{open}}\, C_f > 0$, giving the closed-form limit:

$$V(\lambda) \xrightarrow{\lambda \to 0} \frac{p\, \Delta d}{N_{\text{open}}\, C_f} = \frac{0.008 \times 800}{4 \times 8{,}000} = 2.0 \times 10^{-4} \approx 0.02\%. \tag{B.4}$$

Finite-width Phase I. For any $\varepsilon > 0.02\%$, there exists $\lambda_\varepsilon > 0$ such that $V(\lambda) < \varepsilon$ for all $\lambda \in [0, \lambda_\varepsilon)$. Empirically, $\lambda_\varepsilon \approx 2.00\,\text{trips/km}^2$ for $\varepsilon = 0.5\%$ (Section 6). Combining with Part (i), the interval $[0, \lambda_\varepsilon)$ constitutes a genuine Phase I: $D_{\text{reloc}} > 0$ throughout while $V < \varepsilon$. The limit (B.4) is consistent with the observed $V = +0.41\%$ at $\lambda = 0.20$ (Table 9): the gap between 0.02% and 0.41% reflects stochastic noise over $n = 20$ replications and the non-zero $O(f^5)$ term at finite $\lambda = 0.20$. The non-degeneracy holds for any $\beta > 0$ and positive cost parameters; it is a structural consequence of the $O(f^{\beta+1})$ vs. $O(f^0)$ scaling contrast between BPR congestion savings and fixed infrastructure costs. □

References

Arnott, R., A. de Palma, and R. Lindsey. 1993. A structural model of peak-period congestion. *American Economic Review 83*(1), 161–179.

Chung, S. H., A. W. Siddiqui, and H. L. Ma. 2020. Heterogeneous drone delivery with battery constraints. *Transportation Research Part B 136*, 68–89.

Civil Aviation Administration of China. 2025. 2025 statistical bulletin on civil aviation development. Technical report, Civil Aviation Administration of China (CAAC), Beijing, China.

Cummings, C. and H. Mahmassani. 2024. Airspace congestion, flow relations, and 4-D fundamental diagrams for advanced urban air mobility. *Transportation Research Part C 159*, 104467.

Daganzo, C. F. 1994. The cell transmission model. *Transportation Research Part B 28*(4), 269–287.

Dorling, K., J. Heinrichs, G. G. Messier, and S. Magierowski. 2017. Vehicle routing problems for drone delivery. *IEEE Transactions on Systems, Man, and Cybernetics 47*(1), 70–85.

Kitjacharoenchai, P. and others. 2019. Multiple traveling salesman problem with drones. *Transportation Research Part C 102*, 204–225.

Kopardekar, P. and others. 2016. Unmanned aircraft system traffic management (UTM) concept of operations. In *AIAA Aviation Forum*. AIAA.

Levin, M. W., L. Stuive, and F. Gzara. 2024. Airspace network design for urban UAV traffic management with congestion. *Transportation Research Part C 168*, 104830.

Macrina, G. and others. 2020. Drone-aided routing: A literature review. *Transportation Research Part C 120*, 102762.

Mercan, Tülin, Volkan Yavaş, Dilek Can, and Yasin Mercan. 2025. Vertiport location selection criteria for urban air mobility. *Journal of Air Transport Management 124*, 102700.

Murray, C. C. and A. G. Chu. 2015. The flying sidekick traveling salesman problem. *Transportation Research Part C 54*, 86–109.

Otto, A. and others. 2018. Optimization approaches for civil applications of UAVs. *Networks 72*(4), 411–458.

Revillod, G., J. Kremer, and M. Bierlaire. 2025. Vertiport location and sizing: A methodology for airport operators. *Transportation Research Part E: Logistics and Transportation Review 194*, 103912.

Rimjha, M. and others. 2023. New infrastructures for urban air mobility systems. *Journal of Air Transport Management 111*, 102437.

Roughgarden, T. 2005. *Selfish Routing and the Price of Anarchy*. Cambridge, MA: MIT Press.

Stolaroff, J. K. and others. 2018. Energy use and life cycle greenhouse gas emissions of drones for commercial package delivery. *Nature Communications 9*, 409.

Vickrey, W. S. 1969. Congestion theory and transport investment. *American Economic Review 59*(2), 251–260.

Weng, C. and others. 2025. Urban low-altitude air transport management. *Transportation Research Part C 178*, 105237.