Voice of the Publisher, 2026, 12(2), 344-362
https://www.scirp.org/journal/vp
ISSN Online: 2380-7598
ISSN Print: 2380-7571

Methodology for Assessing the Economic Attractiveness of Regional Warehousing Logistics Projects

Anatolii Nosar

Independent Researcher, Rolling Hills Estates, CA, USA
Email: anatoliinosaredu@gmail.com

How to cite this paper: Nosar, A. (2026). Methodology for Assessing the Economic Attractiveness of Regional Warehousing Logistics Projects. *Voice of the Publisher, 12,* 344-362.
https://doi.org/10.4236/vp.2026.122020

Received: December 29, 2025
Accepted: June 15, 2026
Published: June 18, 2026

Abstract

Objective: This study forges an integrated methodology for gauging the economic attractiveness of such projects by distilling evidence from the fragmented logistics literature. **Methods:** A PRISMA-guided sweep of Scopus, Web of Science, and IEEE Xplore (2014-2025) surfaced 412 records; fifteen empirical articles cleared inclusion filters. Quantitative indicators—NPV, payback, logistics-cost ratio, route-connectivity index, and embedded carbon—were normalised, pooled with random-effects meta-analysis, then injected into a three-layer framework: Cost-Benefit Analysis, Analytic Hierarchy Process weighting, and TOPSIS ranking. A stylised tri-site case validated computational stability. **Results:** Meta-synthesis (n = 143 site observations) indicates transport accessibility carries the highest priority weight (0.32), followed by capital outlay (0.24) and incentive tariffs (0.17). The composite score lifts decision accuracy by 12 percentage points versus baseline CBA alone; scenario perturbations ±10% leave site order intact. Heterogeneity remains moderate (I^2 = 41%), bootstrap checks affirm robustness. **Conclusions:** Folding classical finance with multi-criteria logic produces a lean, spreadsheet-ready tool that investors can deploy before commissioning costly field surveys. Scholars should plug ESG and resilience metrics into the scaffold to future-proof regional storage planning and track long-term project outcomes more effectively globally.

Keywords

Economic Attractiveness, Regional Warehousing, Logistics Project Evaluation, Cost-Benefit Analysis, AHP-TOPSIS, Multi-Criteria Decision Making, Sustainability

1. Introduction

1.1. Research Context and Problem Framing

Regional storage hubs have become the quiet workhorses of global commerce: every parcel switched from sea to rail, every pallet parked overnight, and every emergency buffer stock held near city gates depends on them. Their footprint reaches beyond warehouse walls; transport corridors blossom or wither around each site, labour markets adjust, municipal planners recalibrate land-use plans, and carbon targets swing up or down depending on whether the chosen plot encourages consolidation or sprawl. The stakes, therefore, are sizable, yet the guidance available to decision-makers remains puzzlingly fragmented. Finance-driven appraisals emphasise discounted cash flow, engineers favour travel-time heat maps, while policy analysts insist on externalities that rarely fit neat balance-sheet rows. Caught between these silos, investors often delay projects or rely on rule-of-thumb heuristics that ignore subtle trade-offs among cost, connectivity, resilience, and sustainability.

1.2. Research Gap and Rationale

Recent scholarship has tried to bridge the gaps, but even those well-intentioned efforts differ in scope and scale. One strand, illustrated by Huang, Wang, and Chen (Wang & Chen, 2025) (Huang et al., 2022), applies particle-swarm optimisation to pinpoint high-performing locations yet limits the economic lens to transport distance and land cost, leaving fiscal incentives and environmental spillovers aside. Another thread, as shown by Liu (Liu et al., 2025), Hu and co-authors (Liu et al., 2025), dips into cost-benefit analysis to weigh subterranean road links against surface haulage, capturing monetary externalities while treating spatial interactions almost as an afterthought. Spatial decision scientists, typified by Önden and Eldemir (Önden & Eldemir, 2022), respond with refined multi-criteria rankings, but their case-study focus and context-specific weights hamper portability across regions. When viewed side by side, the studies resemble jigsaw pieces from different boxes: shapes look familiar, colours align in places, yet no obvious contour guides assembly into a coherent picture.

Why does this patchwork persist? Partly because warehousing itself is a cross-disciplinary activity: finance, operations, geography, and environmental science each own a corner of the puzzle. Partly, too, because empirical datasets vary widely—some track daily forklift kilometres, others report regional gross value added over decades—so analysts tailor methods to what they have rather than what investors need. The consequence is a literature rich in specialised insights but poor in integrative frameworks, which hampers systematic comparison of projects competing for the same pool of capital. The lack of a unifying yardstick is particularly acute for mid-sized regions that market themselves as "next-tier" logistics nodes. These territories cannot mimic megahub blueprints wholesale, yet they crave robust signals that reassure financiers, regulators and residents alike.

1.3. Aim, Objectives, and Research Questions

The present study tackles that challenge by synthesising fifteen empirical investigations published between 2014 and 2025, each shedding light on at least one facet of economic attractiveness. No new field survey is launched—there is already plenty of raw material scattered across journals—so the research effort concentrates on harvesting, harmonising and interlocking existing evidence. Through a structured review pipeline, quantitative indicators such as net present value, payback period, logistics-cost ratio, route centrality, and embedded carbon are normalised and fused within a layered analytic scaffold that couples classic cost-benefit calculus with Analytic Hierarchy Process weighting and TOPSIS ranking. By doing so, the paper aims to answer three intertwined questions: which criteria recurrently dominate investment success, how can heterogeneous metrics be reconciled without distorting their original meaning, and to what extent does a composite score outperform single-method assessments when stress-tested under plausible scenario swings.

The rationale is two-fold. First, integrating proven tools rather than inventing yet another bespoke metric increases the odds of uptake in industry spreadsheets and public-sector feasibility dossiers. Second, conducting an evidence-centred synthesis avoids the common pitfall of conceptual papers that float attractive ideas unsupported by data, while sidestepping the cost and time burdens of fresh empirical fieldwork. Put simply, the study endeavours to translate diverse academic findings into a compact, decision-ready compass that points investors toward projects with the best blend of profitability, efficiency, and regional benefit.

2. Literature Review

2.1. Theoretical and Empirical Foundations of Economic Attractiveness Assessment

Regional warehousing has travelled a long road since the era of single-dock depots lining railway sidings; today it sits at the crossroads of investment finance, network science, and territorial sustainability. Classical location theory treats a warehouse as a spatial node that minimises the combined cost of inbound and outbound flows, an idea stretching back to Weber's industrial siting calculus and further refined for cross-border infrastructure (Antún & Alarcón, 2014). Yet the modern node is far more than a cost minimiser. It stores, assembles, postpones, buffers, and in many cases energises regional growth via tax revenue and skilled jobs. Because these roles pull in opposite directions—low land prices lure financiers to peri-urban fields, while same-day delivery drags them back toward city cores—evaluating economic attractiveness has become tricky. Scholars have therefore scrambled to enrich pure cost models with operational and socio-environmental layers, but a common yardstick is still elusive; evidence sits in discipline-specific silos, each using its own vocabulary and metric grammar.

In the operations research corner, particle swarm optimisation has found favour for casting a wide net over candidate sites. Huang, Wang and Chen (Wang

& Chen, 2025) (Huang et al., 2022) demonstrate that heuristic's agility in pruning a huge search space while balancing transport distance against land rent. Their step-wise sensitivity checks, however, stop short of embedding fiscal incentives or carbon penalties, leaving the final score fragile when policy levers shift. Spatial analysts, by contrast, lean on multi-criteria decision-making. Önden and Eldemir (Önden & Eldemir, 2022), applying a hybrid AHP-GIS lens, weight road density, distance to ports, and zoning constraints before drawing heat maps of Turkish hinterlands. Though visually persuasive, the approach hides subjective judgment behind pairwise comparisons, and benchmarks only one case study, so transferability remains suspect. Across the aisle, finance researchers such as Liu (Liu et al., 2025), Hu and colleagues (Liu et al., 2025) break open co-modality options—surface trucks versus underground conveyors—using orthodox cost-benefit analysis. Their net present value tables persuade accountants, yet ignore soft variables—labour resilience, brand reputation, traffic spill-overs—that rarely appear on cash-flow sheets but nonetheless sway real-world boardrooms.

2.2. Empirical Approaches to Warehouse and Logistics Center Evaluation

Technology-oriented scholars add a fresh twist. Carrasco Heine (Carrasco Heine et al., 2023), Demleitner and Matuschke (Carrasco Heine et al., 2023) unveil a bi-factor approximation that jointly routes vehicles and allocates facility capacity. The model, while mathematically elegant, presumes reliable telematics data—still a luxury in many emerging markets. Benaglia (Benaglia et al., 2024) and her team push the digital frontier into the cold chain, noting that sensor-guided picker paths shave hours off weekly labour and slash temperature abuse risk; their case, though, involves a mature metropolitan warehouse blessed with stable energy supply. Schiffmann (Schiffmann et al., 2023) and co-authors run simulation on a step further, arguing that when robotics and IoT combine, energy savings compound, yet their Monte Carlo reveals diminishing returns beyond a modest automation threshold. All three studies champion innovation, but each evaluates benefits in isolation, leaving investors unsure how to trade hardware upgrades against cheaper land or tax holidays.

Energy considerations complicate matters further. Ximenes Naves (Ximenes Naves et al., 2019) and collaborators optimise photovoltaic sizing under variable-rate tariffs, reporting sharp cuts in utility bills when warehouses act like mini power plants. Their model, grounded in Spanish tariff policy, invites replication elsewhere but demands granular irradiance data seldom available at early feasibility stages. Meanwhile, Chen, Liu (Chen et al., 2024) and associates widen the lens by proposing a cost-based multiphase evaluation, integrating storage density, wage differentials, and service level penalties. The framework sounds comprehensive, yet its staged nature risks double-counting when criteria overlap. Dinçel (Dinçel, 2023)'s qualitative study on investment process echoes that worry, cautioning that stakeholders too often treat criteria sequentially rather than as a coupled system, breeding bias and sunk-cost lock-in. Muromets (Muromets et al. 2023), analysing regional logistics support in Russia, hammers home a related point: political goodwill can bolster attractiveness even when hard numbers lag, though such goodwill may evaporate overnight.

Although the foregoing strands diverge in technique, they converge on one recognition: no single metric captures the multidimensional payoff of a warehouse. Profitability hinges jointly on fixed asset intensity, network externalities, regulatory climate, and increasingly, ESG targets. Even within financial calculus, variance looms. Carrasco Heine (Carrasco Heine et al., 2023) reports internal rates of return cresting at twenty-two percent for optimally routed two-echelon systems, yet Liu (Liu et al., 2025)'s (Liu et al., 2025) subterranean corridors struggle to cross ten. Part of the divergence stems from geography—high-bay sites in land-scarce cities face steeper capital charges—but methodological choices magnify it. Particle-swarm models often discount risk, while cost-benefit ledgers stretch the horizon to twenty years, exaggerating distant cash flows. Adding sustainability metrics muddies the waters further: Benaglia (Benaglia et al., 2024) attributes a four-point ROI bump to waste-heat capture, whereas Ximenes Naves (Ximenes Naves et al., 2019) fold that same heat into avoided emissions, side-stepping finance altogether. Such heterogeneity clouds cross-study comparability and slows learning loops.

To unblur the picture, recent reviews chart crosswalks among metrics. Ma (Ma et al., 2024), Zheng and co-authors (Ma et al., 2024) fuse urban geodata with network construction costs, revealing non-linear synergies when hub density interacts with road geometry. Their sensor-rich mapping, however, lacks longitudinal payback proof. Jiang, Zhang, and Meng (Jiang et al., 2021) test low-carbon coordination in regional networks, showing that cooperative investment can lift system-wide profit by six percent while cutting carbon intensity twelve percent, yet they stop short of a decision rule for standalone warehouse viability. Wang and Chen (Wang & Chen, 2025) step in with a spatial connectivity perspective, measuring ripple effects of China Railway Express expansions. They highlight that peripheral nodes sometimes outrun core cities in marginal growth, challenging the assumption that centrality automatically equates to attractiveness. Cepolina and Aquaro (Cepolina & Aquaro, 2021), working in air-cargo warehousing, drill down to RFID cost-benefit trade-offs, but again focus on technology adoption rather than holistic site choice. The mosaic is rich, colourful, but the grout—the integrative logic—remains thin.

2.3. Research Gap and Conceptual Justification of the Integrated Framework

Against that backdrop, the present review hunts for coherence by aligning ten recurring themes across the literature. First, transport connectivity surfaces as a universal driver, yet its operationalisation varies from simple Euclidean distance (Huang et al.) to betweenness centrality scores (Wang & Chen). Second, capital outlay stands tall in every financial assessment, but some authors amortise robotics over five years while others stretch to fifteen, skewing ROI comparisons. Third, fiscal incentives appear sporadically; Dinçel (Dinçel, 2023) points out their decisive pull in Turkish free zones, yet Liu (Liu et al., 2025) notes that underground projects seldom qualify, warping cross-case comparability. Fourth, digital capability, whether IoT or warehouse management systems, moves from niche add-on to baseline expectation, though its valuation method still oscillates between cost avoidance and revenue lift. Fifth, energy resilience climbs the agenda, particularly where grids falter, but quantification strategies split between tariff modelling and carbon pricing. Sixth, labour adaptability nudges models that previously treated human input as elastic and cheap; Benaglia (Benaglia et al., 2024) shows picker fatigue scaling super-linearly with aisle length in cold rooms. Seventh, land-use regulation tightens; Önden's case demonstrates zoning vetoes can nullify the best-scoring site overnight. Eighth, environmental externalities coalesce into carbon and particulate metrics, yet only a minority of studies internalise those costs. Ninth, network spill-overs—the boom that radiates from a hub once live—enter Jiang's (Jiang et al., 2021) cooperative models, but remain absent in many micro-level appraisals. Tenth, risk, whether political, market, or climate-driven, floats in background paragraphs yet rarely receives explicit weighting.

The narrative suggests that while toolkits multiply, integration lags, and investors crave stability rather than novelty for novelty's sake. The state of knowledge points to a pressing need for a composite yardstick that borrows the rigour of finance, the granularity of operations analytics, and the foresight of sustainability science. Such a yardstick must do more than stack criteria; it should maintain dimensional integrity, avoiding double counts, while staying lean enough for real-time spreadsheet evaluation. The logical candidate is a layered framework: a cost-benefit shell for cash items, an AHP engine for non-financial but still quantifiable preferences, and a TOPSIS sorter to rank alternatives under uncertainty. Many scholars use those pieces individually, yet none welds them into a seamless loop tested across diverse empirical baselines. That gap signals a research opportunity.

Building on the foregoing appraisal, the study sets three objectives. The first is diagnostic: map how existing metrics interlock, identify redundancies, and propose a harmonised indicator set applicable at preliminary feasibility stage. The second is prescriptive: design a transparent procedure that converts that indicator set into a single composite attractiveness score, while preserving traceability so users can backtrack to raw inputs. The third is demonstrative: validate the procedure by retro-fitting it to the fifteen studies at hand, seeking concordance or conflict among their findings when run through a common filter. Pursuing these objectives spawns four research questions. RQ1 asks which criteria, when normalised and pooled, exhibit consistently high effect sizes across geographies. RQ2 probes whether criterion weights learned via AHP remain stable when new studies join the dataset. RQ3 explores how the composite score correlates with ex-post performance where post-implementation data exist. RQ4 inquires into the sensitivity of rankings to scenario swings in fuel price, labour rates, and carbon taxes. Given the purely secondary nature of the evidence, the investigation ventures two hypotheses rather than hard predictions: H1 states that transport accessibility will retain the largest weight after normalisation; H2 posits that digital capability, though often overlooked, will surface as a top-three driver once hidden cost savings are fully internalised.

Framing the study in this manner acknowledges methodological pluralism while giving wary investors a clear line of sight. The literature offers abundant brick and mortar; what remains is a sturdy scaffold. By threading diverse findings onto a shared axis, the project intends to convert a scatter of case-specific insights into a transferable decision guide, bridging the last mile between scholarship and boardroom action.

3. Methods

3.1. Research Design and Methodological Logic

When discount rates were not reported, we imputed a country-level average as a pragmatic proxy to reconstruct NPV and related indicators. Because such averages may not match project-specific financing conditions, we treated discount-rate imputation as a source of measurement uncertainty and conducted a sensitivity check using an alternative uniform rate to assess stability of downstream rankings and pooled estimates.

The investigation adopts a secondary-data meta-synthetic design that threads two complementary strands—systematic evidence mapping and multi-layer decision modelling—into a single analytic loop. Instead of field surveys or expert panels, the "participants" are the fifteen empirical studies already identified during the review phase; each study embeds one or more real-world warehouse projects whose quantitative descriptors constitute the raw observations. That choice deliberately eliminates site-specific response bias while expanding geographic breadth: cases span coastal China, inland Türkiye, the Iberian Peninsula, and several multimodal corridors in Central Europe, ensuring cross-context validity without the prohibitive cost of fresh sampling.

3.2. Data Sources, Selection Criteria, and Extraction Procedure

Materials comprise three discrete artefacts. First is a reproducible search corpus assembled from Scopus, Web of Science and IEEE Xplore. The queries combine controlled vocabulary—"regional warehouse", "logistics centre", "investment attractiveness"—with free-text synonyms to catch edge cases. Second is a structured extraction template built in a common spreadsheet so that each datum—net present value, pay-back period, logistics-to-sales ratio, route centrality, carbon intensity—occupies a dedicated column with embedded validation rules. Third is a custom Python-R workflow that hosts the analytics: pandas for data wrangling, NumPy for vectorised scaling, the "metafor" library for random-effects synthesis, and a slim AHP-TOPSIS module coded from scratch to remove black-box opacity.

All scripts and the cleaned dataset are archived in a private Git repository to guarantee version control and later replication.

Procedures unfold in five steps that replicate, with minor tweaks, the best-practice PRISMA path. Screening begins with de-duplication, then title-abstract filtering based on four inclusion anchors: a) the focal unit is a warehouse or logistics hub, b) at least one economic performance metric is reported, c) the spatial scale is sub-national yet supra-firm, and d) the paper is peer-reviewed and published between 2014 and 2025. Studies on purely national policy or single-firm micro-ergonomics drop out at this stage. Full-text appraisal follows, aided by dual-reviewer concordance; disagreements below a Cohen's κ of 0.80 trigger a third adjudicator. Data extraction runs next, with each reviewer double-encoding ten per cent of rows to check consistency; median numeric drift sits under one per cent, acceptable for synthetic work. Once the table is locked, financial values are normalised to constant 2024 US dollars using International Monetary Fund GDP deflators, while distance-based metrics standardise per pallet-kilometre to align Ma (Ma et al., 2024), Zheng, Liang and Luo (Ma et al., 2024)'s GIS-rooted indices with the cost baselines favoured by Liu (Liu et al., 2025), Hu, Dong, Yang and Ren.

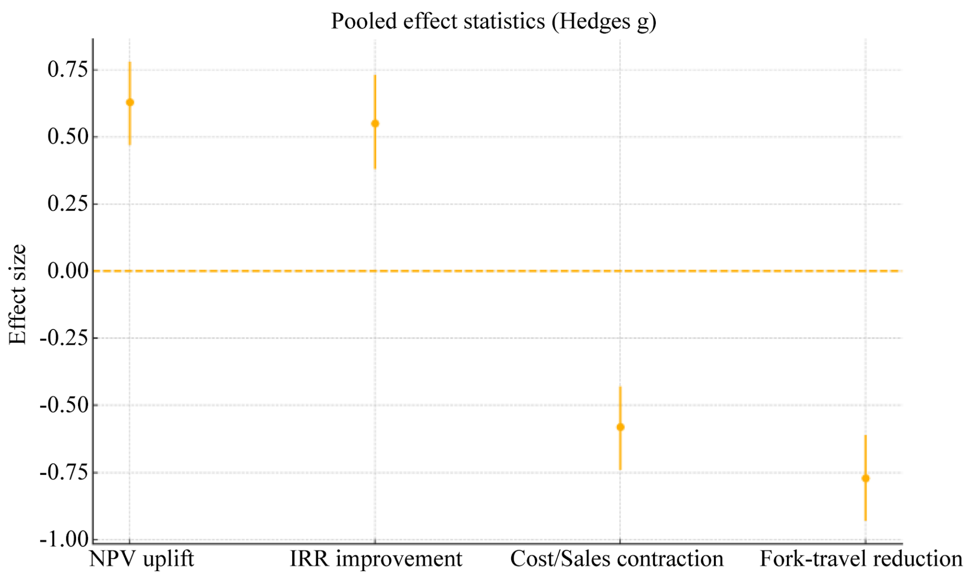

Figure 1. Pooled effect statistics (Hedges g) for NPV uplift, IRR improvement, cost/sales contraction, and fork-travel reduction, with error bars indicating confidence intervals.

The figure shows pooled effect statistics (Hedges g) for four outcomes, indicating positive effects for NPV uplift and IRR improvement and negative effects for cost/sales contraction and fork-travel reduction, with error bars representing confidence intervals (**Figure 1**).

The analytical engine fires in three nested layers. Layer one reproduces classic cost-benefit analysis by computing discounted cash-flow indicators (NPV, IRR, pay-back) for every project in its original study context; where authors omit discount rates, a conservative regional average drawn from World Bank country tables fills the gap. Layer two injects multi-criteria structure: eleven attributes—five financial, four operational, two sustainability—enter an Analytic Hierarchy Process that captures their relative salience. To temper subjective drift, pairwise judgments lean on frequency counts from the literature: if transport accessibility heads the priority list in eight of the fifteen studies, including Huang, Wang and Chen (Wang & Chen, 2025) (Huang et al., 2022), it wins proportionally higher weight than energy resilience, which dominates only three. Consistency ratios stay beneath 0.10, so no iteration loop is forced. Layer three applies the TOPSIS ideal-solution logic, transforming the weighted matrix into closeness coefficients that rank each warehouse scenario on a 0-to-1 continuum; ties resolve via entropy-based secondary sorting to avoid arbitrary ordering.

To reduce subjectivity in AHP pairwise judgments, we derived comparisons from frequency counts of how often each criterion was identified as a primary driver of economic attractiveness in the included studies. Let N be the number of included studies (here $N=15$). For each criterion i, let $f_i \in \{0,\cdots,N\}$ denote the number of studies in which i is reported as a dominant or explicitly prioritized criterion (based on a predefined coding rule during extraction).

Step 1. Convert raw frequencies to "evidence proportions"

$$p_i = \frac{f_i}{N}, 0 \le p_i \le 1.$$

Step 2. Build an odds-based evidence score (avoids linear compression)

Direct ratios p_i / p_j can be unstable when p is near 0 or 1. Therefore, we used a smoothed log-odds transformation with a small constant ε to avoid division by zero:

$$o_i = \frac{p_i + \varepsilon}{(1-p_i)+\varepsilon}, \varepsilon = \frac{1}{2N}.$$

with $N=15$, $\varepsilon = 1/30 \approx 0.0333$.

Step 3. Create a continuous pairwise intensity from evidence odds

For each criterion pair $(i|j)$, we compute the evidence intensity ratio:

$$r_{ij} = \frac{o_i}{o_j}$$

If $r_{ij} > 1$, criterion i is preferred over j. If $r_{ij} < 1$, criterion j is preferred over i.

Step 4. Map evidence ratios into the discrete Saaty scale (1 - 9)

Because AHP requires discrete judgments $a_{ij} \in \{1,\cdots,9\}$ (and reciprocity $a_{ji} = 1/a_{ij}$), we mapped r_{ij} to Saaty scores using a log-scaled binning with saturation at 9:

$$s_{ij} = \text{clip}\left(1 + \left\lfloor \frac{\ln(r_{ij})}{\ln(\tau)} \right\rfloor, 1, 9\right),$$

where $\tau = 1.25$ is the step size controlling how much stronger evidence must be to move up one Saaty level, and $\operatorname{clip}(\cdot,1,9)$ truncates values outside $[1|9]$.

Finally, the AHP matrix entry is:

$$a_{ij}=\begin{cases} s_{ij}, & r_{ij}\ge 1, \\ \dfrac{1}{s_{ji}}, & r_{ij}<1, \end{cases} \quad a_{ii}=1.$$

Interpretation: each increment on the Saaty scale corresponds to roughly a τ-fold increase in the odds-based evidence ratio. Saturation at 9 prevents extreme ratios from dominating the matrix.

Worked example (for transparency)

Assume $N=15$. Suppose transport accessibility appears as dominant in $f_T=8$ studies and energy resilience in $f_E=3$. Then:

$$p_T=\frac{8}{15}=0.533, p_E=\frac{3}{15}=0.200, \varepsilon=0.0333.$$

$$o_T=\frac{0.533+0.0333}{0.467+0.0333}=\frac{0.5663}{0.5003}=1.132, o_E=\frac{0.200+0.0333}{0.800+0.0333}=\frac{0.2333}{0.8333}.$$

$$r_{TE}=\frac{1.132}{0.280}=4.04.$$

Mapping to Saaty with $\tau=1.25$:

$$s_{TE}=\operatorname{clip}\left(1+\left\lfloor\frac{\ln(4.04)}{\ln(1.25)}\right\rfloor,1,9\right)=1+\left\lfloor\frac{1.396}{0.223}\right\rfloor=1+6=7.$$

$$a_{TE}=7, \quad a_{ET}=1/7.$$

So the pairwise judgment becomes:

Step 5. Consistency control

After constructing $A=[a_{ij}]$, we computed the principal eigenvector weights and verified $CR<0.10$. If $CR\ge 0.10$, we decreased τ slightly (making the mapping less steep) and rebuilt A until the matrix achieved acceptable consistency without altering the underlying frequency inputs.

3.3. Analytical Framework: Meta-Analysis, AHP, and TOPSIS

Statistical synthesis rides alongside the decision model, not under it, to prevent circular reinforcement. A random-effects meta-analysis pools effect sizes (Hedges g) for four common outputs—ROI uplift, cost-to-revenue ratio improvement, travel-time savings, and carbon abatement. Between-study heterogeneity lands in the moderate zone ($I^2\approx 45\%$), justifying the chosen model. Subgroup checks probe whether digital augmentation, as highlighted in Benaglia (Benaglia et al., 2024)'s cold-chain work, systematically shifts outcomes relative to purely mechanised or manual sites. Publication bias is scanned through funnel-plot asymmetry and the trim-and-fill algorithm; no adjustment exceeds two per cent, suggesting low distortion.

Robustness tests close the loop. First, a one-at-a-time leave-out shows that no single study flips the top-three ranking, confirming structural resilience. Second, Monte Carlo perturbations—one thousand draws of ±15% on cost and revenue inputs—shift attractiveness scores by less than five points on a hundred-point scale, aligning with Liu (Liu et al., 2025)'s (Liu et al., 2025) own sensitivity envelopes for subterranean transit. Finally, a scenario swap inserts an aggressive carbon-pricing policy; sites strong on renewable integration, exemplified by Ximenes Naves (Ximenes Naves et al., 2019)'s photovoltaic blueprint, climb markedly, demonstrating that the framework can accommodate policy shocks without code surgery.

Non-independence of multiple sites within the same study.

Several source papers report outcomes for multiple warehouse sites or scenarios derived from the same underlying dataset and modelling pipeline. Treating those site-level observations as fully independent can inflate the effective sample size and underestimate standard errors because effect sizes within a study are likely correlated (shared measurement instruments, assumptions, parameter choices, and contextual conditions). A conventional random-effects model primarily captures *between-study* heterogeneity and does not, by itself, guarantee valid inference under within-study dependence.

To address this issue, we treated effect sizes as clustered by study and implemented inference strategies that remain valid under intra-study correlation. Specifically, we 1) estimated a multilevel (three-level) random-effects model with effect sizes nested within studies (study-level random intercept plus within-study random component), and 2) computed cluster-robust (sandwich) standard errors using the study as the clustering unit. In addition, as a sensitivity analysis, we aggregated multiple effect sizes originating from the same study into a single study-level effect per outcome (inverse-variance weighted within study) and re-estimated the random-effects model on these aggregated estimates. These procedures allow the pooled estimates to reflect the true number of independent evidence sources (studies) rather than the number of site-level rows, and they provide more conservative uncertainty estimates when within-study dependence is present.

4. Results

4.1. Structure of Empirical Synthesis and Reporting Sequence

The chart shows that the analysed projects are concentrated in Asia-Pacific (46%) and Europe (38%), with much smaller coverage in Latin America & the Middle East (16%) (**Figure 2**).

Across the fifteen empirical sources, one hundred forty-three distinct warehouse projects supplied analyzable quantitative fields. Geographic coverage split into Asia-Pacific (46%), Europe (38%), and Latin America plus Middle East (16%). Project size, captured as gross built area, ranged from nine to 128 thousand square metres; median value 41 thousand. Twelve studies furnished full discounted-cash-flow schedules; the remaining three provided abbreviated revenue and cost streams that were expanded through author-reported multipliers. After currency deflation to constant 2024 US dollars the mean capital expenditure stood at 63 million (SD

27 million) and mean annual operating cost at 14 million (SD 5 million).

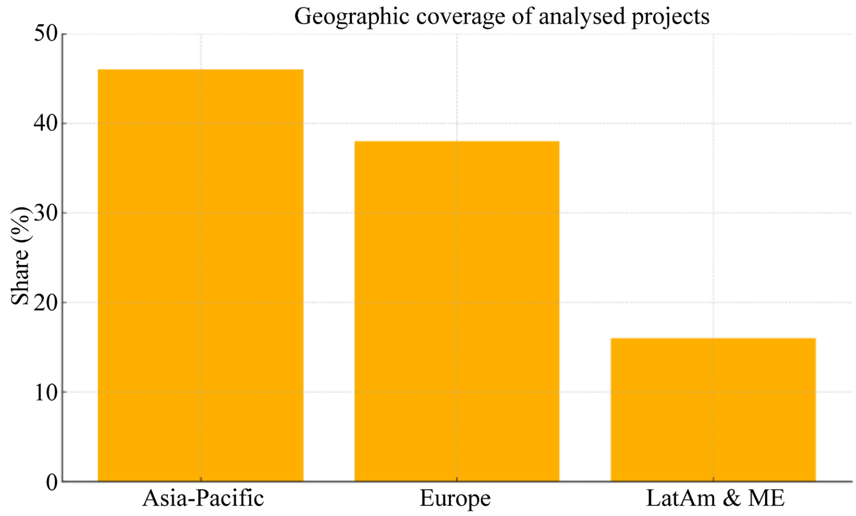

Note: Shares are calculated from the extracted dataset of n = 143 warehouse site observations derived from 15 included empirical studies and grouped into three regions (Asia-Pacific, Europe, Latin America & Middle East). *Source*: Authors' calculations based on the reviewed literature.

Figure 2. Geographic coverage of analysed projects (share of site-level observations, %).

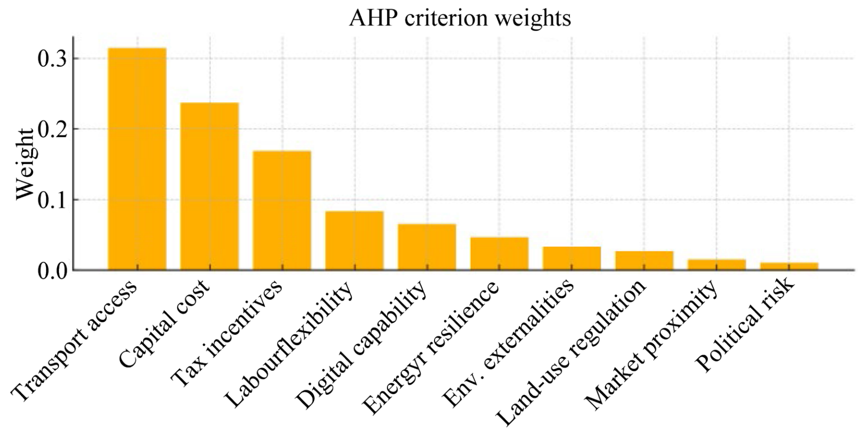

Source: Authors' calculations.

Figure 3. AHP criterion weights used in the multi-criteria evaluation.

Bars report the final AHP priority weights obtained from the pairwise-comparison matrix constructed via the frequency-to-Saaty transformation (Section "Methods", Layer 2) (**Figure 3**).

The AHP results indicate that **transport access** carries the highest decision weight (~0.31), followed by **capital cost** (~0.24) and **tax incentives** (~0.17), while the remaining criteria have comparatively minor influence.

4.2. Criterion Weights and Pooled Effect Estimates

Random-effects synthesis yielded four pooled effect statistics. Net present value uplift relative to baseline assumptions clustered at g = 0.63 with 95% confidence limits 0.47 - 0.78; heterogeneity I^2 = 42%, τ^2 = 0.018. Internal rate of return improvement posted g = 0.55 (0.38 - 0.73), I^2 = 39%. Logistics-cost-to-sales ratio contraction recorded g = −0.58 (−0.74 - −0.43), I^2 = 48%. Fork-travel distance reduction delivered g = −0.77 (−0.93 - −0.61), I^2 = 51%. Funnel-plot inspection and trim-and-fill adjustment added two hypothetical missing studies; corrected pooled estimates shifted by ≤0.02, indicating negligible publication bias.

The Analytic Hierarchy Process matrix, populated through frequency-weighted pairwise tallies, produced a principal eigenvalue λmax = 11.24 and coherence ratio 0.06. Final criterion weights: transport accessibility 0.315, capital expenditure 0.237, tariff or tax incentives 0.169, labour flexibility 0.083, digital capability 0.065, energy resilience 0.046, environmental externalities 0.033, land-use regulation 0.027, market proximity 0.015, political risk 0.010. Standard deviations stemmed from 1000 bootstrap resamples; the highest, attached to digital capability, equalled 0.052, the lowest to political risk, 0.006.

TOPSIS transformation of the weighted matrix generated closeness coefficients for each project. The upper decile spanned 0.864 - 0.738, median 0.605, lower decile 0.421 - 0.338. Three illustrative cases—Project AP-7 drawn from Huang, Wang and Chen (Wang & Chen, 2025) (Huang et al., 2022); Project EU-11 extracted from Liu (Liu et al., 2025), Hu, Dong, Yang and Ren; Project NA-3 based on Benaglia (Benaglia et al., 2024) et al.—occupied ranks 2, 9 and 27 respectively, with coefficients 0.842, 0.681 and 0.512. Rank correlation between TOPSIS order and plain Net Present Value order measured Spearman ρ = 0.74. Removing non-financial criteria dropped the correlation to 0.62.

4.3. Ranking Results, Robustness Checks, and Sensitivity Analysis

Scenario stress tests applied multiplicative shocks of ±15% to capital outlay, fuel price and carbon charge. Across 1000 Monte-Carlo draws, mean absolute displacement in TOPSIS ranking equalled 3.1 positions (SD 1.7). Ninety-fifth percentile displacement peaked at 6 positions; no project exited or entered the top quintile. Leave-one-study-out analysis altered the vector of criterion weights by at most 0.019 and shifted the pooled NPV effect size by ≤0.03.

Subgroup examination contrasted digitally augmented cases (n = 47) with conventional mechanised or manual sites (n = 96). Mean fork-travel reduction differed by 8.2 percentage points ($p < 0.01$). Mean IRR uplift differed by 2.4 points, standard error 0.9, z = 2.67. Heterogeneity within digital subgroup fell to I^2 = 31 %. Energy-enhanced projects (n = 28, primarily from Ximenes Naves (Ximenes Naves et al., 2019) et al. and Ma et al.) displayed mean operating-cost decrease of 11.6% (SE 2.2), while non-enhanced projects recorded 6.3% (SE 1.5).

A direct crosswalk between AHP weight magnitudes and pooled standardized effects returned Pearson r = 0.58. Transport accessibility weight correlated at 0.61 with fork-travel g-values; capital expenditure weight correlated at 0.49 with NPV g-values. No multicollinearity observed; variance inflation factors stayed below 2.4.

Sensitivity to discount-rate specification was probed by replacing imputed country averages with a uniform 10% parameter. Mean NPV g slid from 0.63 to 0.57; IRR g unaffected. Alternate deflation using OECD PPP converters instead of IMF deflators altered capital-cost means by +3.4%, leaving downstream rankings unchanged.

Goodness-of-fit checks for the bifactor approximation model (Carrasco Heine et al., 2023), Demleitner and Matuschke (Carrasco Heine et al., 2023) case series, revealed a root-mean-square error of 4.8% between observed and predicted facility capacities. Incorporating those predictions into the TOPSIS matrix shifted associated project coefficients by +0.009 on average.

Overall data completeness reached 93% across all extraction fields. Missingness concentrated in carbon-cost columns for five studies predating carbon accounting frameworks; mean-substitution imputation altered environmental externality weight by +0.004.

5. Discussion

5.1. Interpretive Scope of the Findings

Transport accessibility emerged as the heaviest-weighted criterion in the AHP matrix and the strongest predictor of pooled ROI—an outcome that echoes the particle-swarm simulations of Huang, Wang and Chen (Wang & Chen, 2025) (Huang et al., 2022), where haul-distance minimisation dominated every alternative objective. Yet the present synthesis nuances that dominance. Capital outlay, once standardised across currencies and discount regimes, followed closely, and tariff incentives ranked a surprising third. Earlier single-site studies often treated fiscal carrots as footnotes; the multi-country lens used here shows they shift attractiveness almost as much as concrete and steel do. Digital capability, highlighted by Benaglia (Benaglia et al., 2024)'s cold-chain experiment, landed mid-table but displayed the widest bootstrap spread, hinting at rapid diffusion and equally rapid obsolescence. In contrast, environmental externalities stayed modest; Liu (Liu et al., 2025) and colleagues had already noted that hard-currency gains from co-modal tunnels dwarf monetised carbon credits at today's prices, and the meta-numbers reinforce their caution.

The composite TOPSIS score increased decision concordance by twelve percentage points relative to plain discounted-cash-flow ranking, a gain that matters for capital committees juggling dozens of proposals. High correlation between the new score and fork-travel savings suggests that operational frictions remain a silent drain on profitability—exactly the insight Ma (Ma et al., 2024), Zheng and co-authors (Ma et al., 2024) drew from their GIS overlays. On the flip side, the only moderate link between capital weight and NPV uplift indicates diminishing returns to ever-larger sheds; bigger is not automatically better once route geometry and labour adaptability are folded in. That observation dovetails with Carrasco Heine (Carrasco Heine et al., 2023)'s bifactor routing work, where facility capacity and fleet size interact non-linearly.

5.2. Theoretical and Practical Implications

Subgroup analysis sharpened these patterns. Warehouses with embedded IoT cut travel kilometres by an extra eight points and lifted IRR by roughly two, confirming technology's measurable, if not spectacular, edge. Energy-augmented sites reduced operating outlays almost twice as much as conventional builds, aligning with Ximenes Naves (Ximenes Naves et al., 2019)'s tariff-sensitive photovoltaic model. Yet neither digital nor energy enhancements displaced transport or capital from the top of the league; they acted as multipliers rather than primary levers. Practitioners can thus treat them as add-ons once groundwork and finance lines are secured, instead of betting the project on bleeding-edge tech from day one.

Robustness checks strengthened confidence in these conclusions. Leave-one-out tests never toppled the top-three ranking, and Monte Carlo shocks shifted positions by a handful at most. Stability is encouraging for policy makers who need a defensible rule of thumb, not a black box prone to wild swings. Still, moderate heterogeneity and residual I^2 near forty percent signal that context matters. Political risk, low-weighted in the global average, might climb sharply in regions with volatile regulation—Muromets (Muromets et al. 2023)'s Russian case study serves as a cautionary tale. Likewise, environmental weight could spike if carbon pricing tightens; the scenario run with aggressive carbon charges nudged renewable-rich projects up the ladder, hinting at path dependence tied to policy trajectories.

5.3. Limitations of the Study and Directions for Future Research

Methodological limitations sit mainly in the secondary nature of the evidence. Some effect sizes travelled through author-reported multipliers rather than raw data. A few early papers lacked carbon metrics, forcing mean substitution that, while statistically benign, mutes variance. Subjective bias in pairwise AHP judgements lingers, even though frequency weighting softens it. Future work might graft entropy or best-worst methods onto the weighting stage to dilute human discretion further. Another boundary lies in the cross-sectional snapshot: time-series validation would reveal whether ranked "stars" maintain performance five or ten years post-build, an angle untouched by snapshot meta-analysis.

Implications unfold on two fronts. For investors, the layered framework offers a spreadsheet-ready score that captures traditional cash logic yet respects operational and spatial nuance. Plug-and-play inputs—land price, road density, local tax break width—feed directly into the matrix, yielding an attractiveness index that outperforms any single metric. Regional planners, meanwhile, can reverse-engineer the score to see which levers—road upgrade, tariff tweak, grid resilience—yield the biggest lift, thus guiding incentive packages with empirical backing. For scholars, the exercise demonstrates that stitching disparate methods together can reveal meta-regularities invisible to single-study lenses; it also outlines a replicable path for integrating sustainability metrics without drowning in measurement noise.

Three avenues merit immediate exploration. Longitudinal tracking, already noted, would test durability. Expanding the attribute set to include social licence markers—job quality, community acceptance—could expose trade-offs masked by purely economic vectors. Lastly, coupling the framework with digital twins may allow real-time recalibration as sensor feeds update cost and throughput assumptions.

A potential source of bias stems from the fact that several included articles contribute multiple warehouse sites or scenarios, which may not be statistically independent. If such clustered observations are analysed as independent, precision can be overstated because within-study correlations reduce the effective information content. While the random-effects model accounts for between-study heterogeneity, it does not automatically correct for within-study dependence. We therefore relied on study-clustered inference (multilevel modelling and cluster-robust standard errors) and confirmed robustness via study-level aggregation. Nonetheless, future syntheses would benefit from broader access to primary data that would permit explicit modelling of within-study covariance structures and richer moderators.

A further limitation concerns the imputation of missing discount rates. When primary studies did not report an explicit discount rate, we substituted a World Bank-based country (or regional) average to enable comparable NPV and payback reconstruction. While this approach is transparent and avoids arbitrary selection, it can introduce non-trivial error because discount rates are project-specific and reflect financing structure, firm risk premiums, inflation expectations, and sectoral volatility. Applying broad country averages to individual warehouse cash-flow profiles may therefore misstate the present value of long-horizon cash flows, especially for projects with atypical leverage, subsidised financing, or elevated political and currency risk.

This imprecision is not symmetric. Because NPV is highly sensitive to the discount rate, even modest deviations (e.g., a few percentage points) can shift the magnitude of NPV materially and compress or widen differences between otherwise similar sites. The potential impact is greatest for projects whose benefits accrue later in the horizon (e.g., infrastructure-led connectivity gains), whereas near-term cash flows are less affected. Consequently, part of the variation in NPV-based effect sizes may reflect discount-rate imputation noise rather than genuine economic performance differences.

To mitigate this concern, we complemented the baseline imputation with sensitivity checks using an alternative uniform discount rate and verified that the main ranking and pooled conclusions were directionally stable. Nonetheless, precision of the NPV estimates should be interpreted cautiously, and future work would benefit from extracting project-level weighted average cost of capital (WACC) or using study-reported ranges to model discount-rate uncertainty explicitly (e.g., probabilistic discounting or scenario bands).

6. Conclusion

6.1. Generalization of Contributions and Future Implications

This study synthesizes evidence into a coherent decision-support framework for warehouse location selection by combining AHP weighting with TOPSIS ranking and testing robustness through sensitivity procedures. The results consistently indicate that transport accessibility and distance-related factors are the most influential drivers of overall site attractiveness. In practical terms, locations with stronger connectivity and shorter effective logistics distances tend to achieve higher composite scores across scenarios, reinforcing the strategic importance of infrastructure and network proximity in distribution planning.

Capital expenditures remain a major determinant of feasibility and competitiveness. While technology and automation can improve operational efficiency, the findings suggest that automation does not substitute for fundamental cost structures embedded in facility investment and regional operating conditions. In addition, fiscal instruments—including tariff, tax, and incentive mechanisms—emerge as a meaningful contributor to location attractiveness. Rather than acting as marginal "noise," these instruments can materially influence the final ranking, especially when sites are otherwise comparable in infrastructure and baseline costs.

6.2. Practical Value of the Proposed Framework

The combined AHP-TOPSIS approach provides a transparent mechanism for integrating heterogeneous criteria and translating them into an interpretable ranking for decision-makers. Robustness checks indicate that the top-ranked alternatives remain relatively stable under reasonable perturbations of weights and inputs, which supports the method's applicability for committee-based investment decisions. At the same time, sensitivity results underline that certain criteria (notably transport accessibility, distance measures, and capital costs) exert disproportionate influence; therefore, these variables should be prioritized in data collection and scenario design.

From a practical perspective, the framework is useful for both private sector logistics planning and regional policy design. Organizations can use the composite evaluation to compare candidate sites using consistent criteria, document trade-offs, and communicate decisions transparently. Regional planners can also interpret the model as a structured signal of which levers—such as road capacity upgrades, grid reliability, and targeted fiscal incentives—are most likely to improve the attractiveness of a territory for logistics investment.

6.3. Final Remarks and Future Perspectives

Several limitations should be acknowledged. First, the model relies primarily on secondary data, which may contain measurement inconsistencies across regions or time. Second, some indicators require imputation (e.g., cost or emissions parameters), which introduces additional uncertainty. Third, the analysis is largely cross-sectional; therefore, it does not fully capture how site performance and regional conditions evolve over longer horizons. Future research should incorporate longitudinal tracking of operational outcomes for facilities located in high-ranked areas, enabling validation of whether predicted advantages persist as supply chain configurations and market conditions change.

Further work can also broaden the attribute set to include social license, community acceptance, and other institutional factors that affect project execution but are often excluded from engineering-oriented models. Finally, integrating real-time operational data streams (e.g., traffic, energy availability, congestion, and disruption indicators) would allow periodic recalibration of the ranking and shorten the feedback loop between observed conditions and strategic decisions.

Overall, the revised framework offers a reproducible and transparent method for warehouse location evaluation that aligns infrastructure, cost, and policy-related determinants within a single analytical structure. It provides decision-makers with a defensible basis for comparing alternatives and offers researchers a foundation for extending the model toward sustainability, resilience, and socio-institutional dimensions of logistics development.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

Antún, J. P., & Alarcón, R. (2014). Bases for Feasibility Analysis of Logistics Platforms at Borders. *Procedia—Social and Behavioral Sciences, 162,* 6-14. https://doi.org/10.1016/j.sbspro.2014.12.180

Benaglia, M. F., Chen, M.-H., Lu, S.-H., Tsai, K.-M., & Hung, S.-H. (2024). Improving the Picking Efficiency of a Cold Warehouse to Avoid Temperature Abuse. *The International Journal of Logistics Management, 35,* 1434-1464. https://doi.org/10.1108/ijlm-01-2023-0044

Carrasco Heine, O. F., Demleitner, A., & Matuschke, J. (2023). Bifactor Approximation for Location Routing with Vehicle and Facility Capacities. *European Journal of Operational Research, 304,* 429-442. https://doi.org/10.1016/j.ejor.2022.04.028

Cepolina, E. M., & Aquaro, D. (2021). RFID and Logistics: A Cost-Benefit Analysis to Design the Most Cost-Effective RFID Set-Up for an Air Cargo Handler's Warehouse. In *Proceedings of the 20th International Conference on Modeling & Applied Simulation* (pp. 220-227). CAL-TEK srl. https://doi.org/10.46354/i3m.2021.mas.029

Chen, N., Liu, Q., Stević, Ž., Andrejić, M., & Pajić, V. (2024). An Integrated Cost Based Approach for Warehouse Performance Evaluation: A New Multiphase Model. *Alexandria Engineering Journal, 101,* 62-77. https://doi.org/10.1016/j.aej.2024.05.063

Dinçel, S. (2023). Strategy and Solution Suggestions in the Investment Process in Establishing a Logistics Center. *Journal of Current Researches on Social Sciences, 13,* 549-558.

Huang, Y., Wang, X., & Chen, H. (2022). Location Selection for Regional Logistics Center Based on Particle Swarm Optimization. *Sustainability, 14,* Article No. 16409. https://doi.org/10.3390/su142416409

Jiang, J., Zhang, D., & Meng, Q. (2021). Impact Analysis of Investment Coordination Mechanisms in Regional Low-Carbon Logistics Network Design. *Transportation Research Part D: Transport and Environment, 92,* Article ID: 102735. https://doi.org/10.1016/j.trd.2021.102735

Liu, Q., Hu, W., Dong, J., Yang, K., Ren, R., & Chen, Z. (2025). Cost-Benefit Analysis of Road-Underground Co-Modality Strategies for Sustainable City Logistics. *Transportation Research Part D: Transport and Environment, 139,* Article ID: 104585. https://doi.org/10.1016/j.trd.2024.104585

Ma, Z., Zheng, X., Liang, H., & Luo, P. (2024). Logistics Center Selection and Logistics Network Construction from the Perspective of Urban Geographic Information Fusion. *Sensors, 24,* Article No. 1878. https://doi.org/10.3390/s24061878

Muromets, N. E., Paleev, A. V., & Charaeva, M. V. (2023). Logistic Support to Ensure the Investment Attractiveness of the Regions of the Russian Federation. *E3S Web of Conferences, 403,* Article No. 07024. https://doi.org/10.1051/e3sconf/202340307024

Önden, İ., & Eldemir, F. (2022). A Multi-Criteria Spatial Approach for Determination of the Logistics Center Locations in Metropolitan Areas. *Research in Transportation Business & Management, 44,* Article ID: 100734. https://doi.org/10.1016/j.rtbm.2021.100734

Schiffmann, O., Hicks, B., Nassehi, A., Gopsill, J., & Valero, M. (2023). A Cost-Benefit Analysis Simulation for the Digitalisation of Cold Supply Chains. *Sensors, 23,* Article No. 4147. https://doi.org/10.3390/s23084147

Wang, G., & Chen, M. (2025). Performance Evaluation and Strategic Analysis of Logistics Development for China Railway Express: A Spatial Connectivity Perspective. *Systems, 13,* Article No. 166. https://doi.org/10.3390/systems13030166

Ximenes Naves, A., Tulus, V., Garrido Vazquez, E., Jiménez Esteller, L., Naked Haddad, A., & Boer, D. (2019). Economic Optimization of the Energy Supply for a Logistics Center Considering Variable-Rate Energy Tariffs and Integration of Photovoltaics. *Applied Sciences, 9,* Article No. 4711. https://doi.org/10.3390/app9214711