International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 10 | Oct - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

E-Commerce Delivery Optimization: A Transportation-Problem Perspective on Fulfilment-to-Hub Allocation and Last-Mile Routing

Aarush Gangoli, Aashna Bhansali, Akshat Kala, Amoolya Kandra, Anagh Gupte

Anil Surendra Modi School of Commerce (ASMSOC), SVKM'S NMIMS, Mumbai, Maharashtra, India

Corresponding author: Amoolya Kandra, amoolyakandra@gmail.com

Course context: Operations Research, B.Sc. Finance, Semester III

Article type: Review and applied modelling study

Abstract

Background. As online retailing has grown into an essential part of the global economy, delivery of goods has become one of the largest expenditures of such firms, with the last-mile delivery from the hub to the customer being the most expensive one. Operations research (OR) provides ways to optimize both the transportation of goods from the company to customers and their movement within supply chains. Problems range from the classical transportation problem (TP) to vehicle routing.

Objective. This paper (i) provides a discussion of two journal articles in which OR is used to optimize delivery in the context of e-commerce, (ii) demonstrates that TP is embedded in the studies as the structure of a routing model and order consolidation, (iii) outlines five research gaps that were not addressed by the authors, and (iv) extend the TP framework to address each of the gaps with numerical experiments.

Methods. The articles by Cokyasar et al. (2023), which addresses a time-constrained capacitated vehicle routing problem for four US cities, and by Muñoz-Villamizar, Velázquez-Martínez and Caballero-Caballero (2024), which describes a large-scale consolidation model for a Mexican retailer, are analyzed. A TP with three fulfillment centers and four hubs is solved with the North-West Corner, Least Cost and Vogel's starting allocations, with results verified with the MODI method and linear programming. Five extended TP models are tested: a pooled multi-carrier TP, time-indexed fixed-charge TP for consolidation, two-stage stochastic TP for peak season, bi-objective cost–emissions TP, and a TP with returns to base.

Results. The base case of TP results in an optimal plan of INR 40,400 per day, which is 9.4% better than the one produced with the North-West Corner method. Pooled allocation of two carriers' depots reduces their line-haul cost by 11.7% and parcel-kilometers by 18.3%. A one-day delivery deadline makes it possible to reduce the number of vans by 19.0% after accounting for the extra cost of customer waiting time. Booking capacity from multiple carriers based on scenarios with demand equal to the peak, as opposed to average daily demand, makes it possible to reduce delivery costs by INR 4,651 per day (value of the stochastic solution). Emissions from electric or CNG vans can reduce the carbon footprint by 30% with a 7.4% increase in costs, while returned parcels add 71% to the cost per delivered parcel and make the base case of TP infeasible.

Conclusion. The cost of delivery in e-commerce is determined at three stages: allocation of packages to carriers, scheduling of their pickup and drop-off with customers, and routing of the vehicles between facilities. The TP model provides a transparent and computationally inexpensive way to analyze the first two stages, with the third one addressing the questions that routing-specific research has not resolved. A sequential approach in which allocation, consolidation, and routing are optimized in that order has theoretical and practical value.

Keywords: e-commerce logistics; transportation problem; Hitchcock model; vehicle routing; last-mile delivery; order consolidation; collaborative logistics; stochastic programming; operations research

1. Introduction

Online retail requires rapid and cheap transportation of enormous quantities of small parcels. Such a typical order passes from a seller or warehouse to a fulfilment centre (FC), line-hauled to a city delivery hub and finally delivered to the customer's door. Each step presents a decision on how much to ship from where to where, and the last step adds an ordering decision over the customers. Both are classical OR problems.

The first being the transportation problem (TP), formulated by Hitchcock (1941), studied independently by Koopmans (1949) and given an efficient solution method by Dantzig's (1951) specialisation of the simplex method. The second is the vehicle routing problem (VRP), introduced as the truck dispatching problem by Dantzig and Ramser (1959) and extended in many directions since (Toth & Vigo, 2014). Since the TP is solvable polynomially but the VRP is NP-hard, practitioners tend to use them in sequence, solving the allocation exactly and routing within each hub's catchment.

Recent OR literature on e-commerce delivery has focused on the routing end of this chain: drones, lockers, crowd-shipping, electric vans and same-day dispatch (Boysen, Fedtke & Schwerdfeger, 2021; Savelsbergh & Van Woensel, 2016), with the allocation layer either as solved or fixed. This paper argues that the transportation model is still useful for e-commerce, because (1) it is already present inside the routing and consolidation models that the literature uses, as the flow-balance structure that moves parcels from a depot or a day of arrival to a point of delivery, and (2) its simple structure makes it easy to extend in exactly the directions that the routing studies leave open: sharing capacity between carriers, choosing when to dispatch, planning under uncertain demand, trading cost against emissions, and handling failed deliveries.

Research questions.

- **RQ1.** What real-life delivery problems in e-commerce are addressed by OR, and why are OR techniques suitable?
- **RQ2.** How are the objectives, constraints and models structured in the reviewed studies, and where does the transportation problem appear inside them?
- **RQ3.** What do the solution methods and results imply for practice, and what are their strengths and limits?
- **RQ4.** Which questions do the reviewed studies leave unanswered, and how far can extensions of the transportation model answer them?

Contributions. (1) A structured, detailed comparison of two e-commerce delivery optimisation studies. (2) A unified view that shows the transportation structure inside a CVRP and inside a consolidation model. (3) A complete worked solution of the base transportation model with all three classical starting rules, a MODI optimality test, shadow prices and cost ranging. (4) A gap matrix covering five gaps in the reviewed work. (5) Five extended transportation models, each solved and interpreted, that address those gaps.

Section 2 describes article selection and Section 3 the real-life context. Section 4 reviews the two articles in depth and Section 5 sets out the literature gap. Section 6 presents the models, Section 7 the solution methods, Section 8 solves the base model, Section 9 reports the gap-filling analysis, Section 10 discusses the results and Section 11 concludes.

2. Selection of Research Articles

Two articles were selected because they (a) focus on e-commerce parcel delivery, (b) use explicit mathematical programming, (c) report results on numerical data using real or realistic data and (d) differ in terms of scale, geography and method, allowing a useful comparison. The two articles were published in Transportation Research Record and Expert Systems with Applications, respectively, and are both peer-reviewed journals in the transportation and applied-optimisation literature.

Table 1. Reviewed articles.

	Article A	Article B
Citation	Cokyasar, T., Subramanyam, A., Larson, J., Stinson, M., & Sahin, O. (2023). Time-constrained capacitated vehicle routing problem in urban e-commerce delivery. Transportation Research Record, 2677(2), 190–203.	Muñoz-Villamizar, A., Velázquez-Martínez, J. C., & Caballero-Caballero, S. (2024). A large-scale last-mile consolidation model for e-commerce home delivery. Expert Systems with Applications, 235, 121200.
Domain	Urban parcel delivery by conventional and battery-electric trucks	Home delivery with flexible (longer) delivery windows
Problem type	Capacitated VRP with maximum route time (and distance for BEVs)	Order consolidation over days plus routing
Technique	Exact mixed-integer programme solved with Gurobi, benchmarked against an iterated tabu search	MILP plus a custom metaheuristic for large instances
Setting	Austin, Bloomington, Chicago and Detroit (U.S.), demand from the POLARIS agent-based simulation	Data from a large Mexican retailer
Headline result	Joint operation of depots cuts vehicle miles travelled by 38.8% in Austin	Cost savings of up to 52%, with fewer vehicles needed

3. Real-Life Context

3.1 The problem

E-commerce growth has resulted in an increase in last-mile transportation, often home delivery, with customers demanding fast transit. The final leg is usually the costliest per parcel since vehicles carry small drops to scattered addresses, with frequent first-attempt failures (Boysen et al., 2021). Small percentage gains thus compound across millions of parcels. The problem is three-tiered: where each parcel is sourced from (allocation), when they are sent (dispatch timing) and how vehicles visit customers (routing).

3.2 Article A: urban parcel trucks under capacity and time limits

Cokyasar et al. (2023) address delivery trucks that depart a distribution centre, serve customers within a maximum allowed time and return. Since vehicle capacity and the time limit already restrict the route length, mid-route recharging is unnecessary for battery-electric vehicles (BEVs). This allows the authors to compare conventional vehicles and BEVs within one model. They study the impact of capacity, maximum travel time, service (dwell) time and BEV range on system measures such as vehicle miles travelled (VMT), fleet size and energy use. The demand comes from POLARIS, an agent-based travel simulation developed at Argonne National Laboratory, and covers one simulated day of household orders served by the depots of several providers (Amazon, FedEx, UPS and USPS).

3.3 Article B: consolidation of delivery orders

Muñoz-Villamizar et al. (2024) observe that current systems serve orders first-in-first-out (FIFO) and minimize distance through an ordinary VRP, even when customers accept longer delivery windows, sometimes in return for an incentive from the retailer. They propose delivering orders in consolidated batches across days, so that vehicles leave fuller. The model handles daily time windows and anticipates future orders, and the authors test it with data from a major Mexican retailer.

3.4 Why OR techniques are suitable

- **Clear economic objective.** Cost, distance or time is measurable and additive.
- **Hard physical constraints.** Vehicle capacity, driver hours and FC capacity are naturally linear or integer constraints.
- **Large combinatorial choice.** The number of possible routes, allocations or dispatch schedules grows explosively, so intuition leaves money on the table.
- **Repeatability.** Plans are re-solved daily, so a model amortises its development cost.
- **What-if capability.** Models quantify a policy change, such as a longer delivery window or depot sharing, before it is tried in operations.

4. In-Depth Review of the Two Articles

4.1 Article A: data, model and findings

Data. The study includes four metropolitan areas of highly varying size: daily orders range from 2,816 in Bloomington to 606,669 in Chicago, served from 22 to 53 depots per city. Street-level data are unavailable in the mesoscopic simulation; consequently, customers are aggregated into "super-locations", and customer-to-customer travel time estimated as Manhattan distance divided by an assumed speed of 15 mph.

Model. The formulation is a single-commodity flow model on a complete graph of depot and customer nodes. Binary variables x_{ij} indicate whether a truck uses arc (i, j); continuous variables y_{ij} accumulate the number of packages on board along that arc; further continuous variables track cumulative travel time, and for BEVs distance; an integer variable counts vehicles. The objective is to minimize total distance travelled. Constraints enforce conservation of vehicle flow at every node, each customer is visited only once, the package flow is bounded by capacity ($y_{ij} \leq Q \cdot x_{ij}$), and route time (and BEV distance) in the same way.

Where the transportation problem sits. The variables counting packages on arc ij form a transportation network in their own right: the depot is a supply node, all customers are demand nodes, and the flow balance at each node is exactly the supply and demand constraint of the Hitchcock model. What makes this a hard routing problem is the coupling $y_{ij} \leq Q \cdot x_{ij}$, i.e. parcels can only travel on arcs that a vehicle actually drives. If that coupling is removed, the model reduces to a minimum-cost flow problem, which is solvable in polynomial time. This reading explains why the paper's most striking result, the 38.8% VMT saving from joint depot operation in Austin, is partly an allocation effect: when providers share depots, each customer can be supplied from the nearest depot of any provider, a change to the supply side of the transportation network, before any routing takes place.

Method and computation. The MIP is solved with Gurobi 9.1 on a 40-core server, and compared to an iterated tabu search (ITS) from the literature. The authors report that ITS is faster in every scenario, while the exact model provides a quality benchmark, noting that the exact solution becomes challenging beyond about 100 super-locations.

Findings. Service time per stop is the most important factor: increasing dwell time from zero to five minutes per customer increases VMT by between 10.9% and 33.1% across cities. Vehicle capacity is second: doubling capacity from 120 to 240 packages reduces the required fleet by 35% in Austin, 33% in Chicago and 27% in Detroit. BEV range has little effect, since the average truck travels about 72 miles a day, below the assumed 80-mile range. Finally, joint operation of depots by competing providers reduces VMT by 38.8% in Austin.

4.2 Article B: data, model and findings

Data. The model is tested with data from a major Mexican retailer serving home-delivery customers. Customers may accept a later delivery in exchange for an incentive, which gives the planner a window of several days for each order.

Model. The problem of consolidation is formulated as a mixed-integer linear programme. Each order has a release day and a latest delivery day. The model decides on which day within this window each order is dispatched and on which vehicle it is carried, so that vehicles leave fuller. The objective is total transportation cost (distance, time and vehicle use); the constraints are vehicle capacity, window feasibility and a single delivery of each order. The model also anticipates orders that will arrive on later days.

Where the transportation problem sits. The dispatch-timing element of the model is a transportation problem over time. Each release day is a supply node having the orders that arrived on it; each dispatch day is a demand node with a capacity equal to the number of vehicles sent; and an arc from release day *t* to dispatch day *s* exists only if *s* lies within the order's window. This is the same structure as the classical production-scheduling transportation problem found in OR textbooks (Taha, 2017), but with a fixed cost per vehicle added. Section 9.2 builds this model explicitly.

Method and findings. As the MILP is too large to solve at realistic scale, the authors design a custom metaheuristic. Consolidation increases vehicle utilisation, and so, compared with FIFO dispatch, saves on distance and time, reduces the number of vehicles required, and yields cost savings of up to 52%.

4.3 Side-by-side comparison

Table 2. Comparison of the reviewed studies.

Dimension	Article A	Article B
Decision level	Operational (one day of routes)	Tactical–operational (multi-day dispatch)
Objective	Minimise total distance (VMT)	Minimise total transport cost
Key decision variables	Arc use, package flow, time, distance, fleet size	Dispatch day and vehicle for each order
Key constraints	Capacity, route time, BEV range, visit once	Capacity, delivery windows, deliver once
Embedded TP structure	Package flow from depot to customers	Order flow from release day to dispatch day
Solution method	Exact MIP (Gurobi) vs. iterated tabu search	MILP plus custom metaheuristic
Data	Simulated demand, four U.S. cities	Retailer data, Mexico
Uncertainty	Deterministic	Anticipates future orders
Environmental metric	Energy use, CV vs. BEV	Not central
Main lever for savings	Capacity and depot sharing	Longer delivery windows

5. Literature Gap

The comparison between the two studies under analysis and the broader surveys of the last-mile OR (Boysen et al., 2021; Savelsbergh & Van Woensel, 2016) highlights five questions that both studies fail to address. Each of the questions has some bearing on practical implementation, and it can be answered with recourse to an extension of the transportation model, which is transparent enough to be used in a course project, but yields practically relevant results.

Table 3. Gap matrix (✓ addressed, ◐ partly addressed, ✗ not addressed).

Gap	Article A	Article B	This paper
G1. Allocation of demand across depots, and how sharing changes it	◐ (one case study, routing only)	✗	✓ Section 9.1
G2. The trade-off between consolidation savings and the cost of making customers wait	✗	◐ (assumes customers accept)	✓ Section 9.2

Gap	Article A	Article B	This paper
G3. Uncertain demand, especially sale and festive peaks	✗ (deterministic)	◐ (anticipates orders)	✓ Section 9.3
G4. Cost versus emissions as an explicit trade-off	◐ (energy reported, not optimised)	✗	✓ Section 9.4
G5. Failed deliveries and returns, typical of cash-on-delivery markets such as India	✗	✗	✓ Section 9.5

G1. Allocation and sharing. Article A shows that the joint depot operation saves 38.8% of VMT, but reports that for one city and inside a large routing model, so the reader cannot see how much of the saving comes from re-allocating customers to nearer depots and how much from improved routes. Nor does it ask how the savings should be split between competing providers, which is a precondition for any real collaboration (Guajardo & Rönnqvist, 2016; Cleophas et al., 2019).

G2. The price of waiting. Article B assumes that the customer accepts a longer window, possibly with an incentive, but the value of that incentive is not modelled as a cost inside the optimisation. A planner needs to know how much consolidation is worth once delay has a price, and how long a window must be before extra days stop paying.

G3. Demand uncertainty. Both studies plan for a known set of orders. E-commerce demand is highly seasonal, with sale events and festive peaks, and capacity decisions such as staff an FC or booking trucks must be made before demand is known. Two-stage stochastic programming (Birge & Louveaux, 2011) is the standard way to value such decisions, but is absent from both studies.

G4. Emissions as an objective. Article A compares conventional vehicles and BEVs and reports energy use, but cost and emissions are never traded off in one model. A Pareto frontier would show how much cleaner the network can be made and at what marginal cost per kilogram of CO_2.

G5. Failed deliveries and returns. Neither study models parcels that are not delivered and travel back. In markets where cash on delivery is common, such as India, a share of parcels is refused at the door and returned to origin (RTO). These parcels consume capacity twice and are concentrated in particular catchments, and can change the optimal allocation and not only its cost.

6. Modelling and Problem Formulation

6.1 Classification of the problems

Table 4. Layers of the delivery chain.

Layer	Decision	Model class	Complexity
FC to hub	Parcels shipped per lane	Transportation problem (LP)	Polynomial
Release day to dispatch day	When each order leaves	Time-indexed TP with fixed charges (MILP)	NP-hard in general
Hub to customers	Which truck serves whom, in what order	Capacitated VRP with time limit (MIP)	NP-hard

6.2 The base transportation model (allocation layer)

Sets and indices. FCs $i = 1, \ldots, m$; delivery hubs $j = 1, \ldots, n$. **Parameters.** s_i = daily parcel capacity of FC i; d_j = daily parcel demand of hub j; c_{ij} = line-haul cost per parcel on lane (i, j). **Decision variable.** $x_{ij} \geq 0$ = parcels shipped from i to j per day.

$$\text{Minimise } Z = \Sigma_i \Sigma_j c_{ij} x_{ij} \quad (1)$$

$$\Sigma_j x_{ij} \le s_i \text{ for every FC } i \text{ (supply)} \quad (2)$$

$$\Sigma_i x_{ij} = d_j \text{ for every hub } j \text{ (demand)} \quad (3)$$

$$x_{ij} \ge 0 \quad (4)$$

Assumptions. Linear costs per parcel; one product class; deterministic demand; no capacity constraints on lanes; divisible shipments. Because the constraint matrix is totally unimodular, integer values for supplies and demands yield integer optimal flows. If the total supply exceeds total demand, a dummy hub with zero cost is created to take the excess, and then the problem is balanced. The dual of the model defines a price u_i for each FC and v_j for each hub such that $u_i + v_j \le c_{ij}$, with equality on every lane used; these prices are exploited in the MODI test of Section 8.

6.3 Capacitated VRP with maximum route time (Article A class)

A compact flow-based statement of the problem class considered by Cokyasar et al. is given below. Nodes are $N = \{0\} \cup C$, where 0 is the depot and C the customers. Parameters: q_a = parcels for customer a; Q = vehicle capacity; t_{ab} = travel time; τ_a = service time; T = maximum route time; D_{ab} = distance. Variables: $x_{ab} \in \{0,1\}$ (arc used), $y_{ab} \ge 0$ (parcels on board along the arc), $z_{ab} \ge 0$ (elapsed time on arrival), k (number of vehicles).

$$\text{Minimise } \Sigma_{(a,b)} D_{ab} x_{ab} \quad (5)$$

$$\Sigma_b x_{ab} = \Sigma_b x_{ba} = 1 \text{ for every customer } a;\ \Sigma_b x_{0b} = k \quad (6)$$

$$\Sigma_b y_{ba} - \Sigma_b y_{ab} = q_a \text{ for every customer } a \text{ (transportation-type flow balance)} \quad (7)$$

$$y_{ab} \le Q x_{ab};\ z_{ab} \le T x_{ab};\ \text{time propagation with } t_{ab} \text{ and } \tau_a \quad (8)$$

Constraint (7) is the demand constraint of a transportation problem applied at every customer. Constraint (8) is what makes the problem hard. Article A's modelling insight is that, once Q and T bind, route length is limited and BEV range need not be a separate constraint unless it is shorter than the distance implied by T.

6.4 Consolidation model (Article B class)

The problem of Article B features a batching layer above routing. Orders have their release days and latest delivery days. The model picks a dispatch day for each order and assigns the orders to vehicles so that vehicles can depart fuller to reduce the transportation cost under the vehicle capacity, window feasibility and single delivery constraints. An explicit transportation-style version of its timing layer is given in Section 9.2 with a price on delay.

6.5 Extended transportation models built in this paper

Each extension changes one element of the base model (1)–(4) so that the effect of each gap can be seen in isolation.

- **E1, pooled multi-carrier TP (G1).** Each carrier p has its own depots and customers. The siloed model solves one TP per carrier; the pooled model solves a single TP in which any depot may serve any carrier's demand. The saving is the difference in optimal cost, and since the siloed plan is feasible for the pooled model, the pooled cost can never be higher.
- **E2, time-indexed fixed-charge TP (G2).** Supply node t holds the a_t parcels released on day t; demand node s is dispatch day s. Variables x_{ts} (parcels released on t and sent on s, allowed for $t \le s \le t + w$) and integer n_s (vans sent on s). Minimise $\Sigma_s F \cdot n_s + \Sigma_{t,s} (c + h \cdot (s - t)) \cdot x_{ts}$ subject to $\Sigma_s x_{ts} = a_t$ and $\Sigma_t x_{ts} \le Q \cdot n_s$, where F is the fixed cost per van, c the

variable cost per parcel, h the delay penalty per parcel-day and w the window length. With $w = 0$ the model reproduces FIFO dispatch.

- **E3, two-stage stochastic TP (G3).** First stage: book capacity $b_i \le S_i^{\max}$ at FC i at cost k_i per parcel before demand is known. Second stage, for each scenario ω with probability π_ω: ship $x_{ij\omega} \le b_i$ in total from FC i and cover any unmet demand $u_{j\omega}$ with a third-party (3PL) carrier at penalty p. Minimise $\Sigma_i k_i b_i + \Sigma_\omega \pi_\omega[\Sigma c_{ij}x_{ij\omega} + p \Sigma_j u_{j\omega}]$.
- **E4, bi-objective TP with mode choice (G4).** Each lane can be served by diesel (cost c_{ij}, emissions $e^D \cdot \mathrm{km}_{ij}$ per parcel) or by electric/CNG vans (cost $1.18 \cdot c_{ij}$, emissions $e^E \cdot \mathrm{km}_{ij}$). The ε-constraint method (Haimes, Lasdon & Wismer, 1971) minimises cost subject to total emissions ≤ ε, for a range of ε.
- **E5, RTO-adjusted TP (G5).** If a share r_j of parcels sent to hub j fail and are re-sent, delivering d_j parcels requires shipping $d_j/(1 - r_j)$, and each shipment carries an expected return leg costing $\rho \cdot r_j \cdot c_{ij}$. The model uses effective costs $c_{ij}(1 + \rho\, r_j)$ and effective demand $d_j/(1 - r_j)$, and adds a 3PL source at cost p when FC capacity is short.

Each extension changes one element of the base model (1)–(4) so that the effect of each gap can be seen in isolation.

Table 5. Comparison of the models.

Feature	Base TP	Article A	Article B	Extensions E1–E5
Objective	Min shipping cost	Min distance (VMT)	Min transport cost	Cost; cost + delay; expected cost; cost vs. CO_2
Variables	Continuous flows	Binary arcs, flows, times	Binary assignment and timing	Flows plus integer vans (E2), booking (E3), mode (E4)
Key constraints	Supply, demand	Capacity, route time	Capacity, windows	As base, plus scenario, window, emissions limits
Exact at scale?	Yes	Moderate instances	Needs metaheuristic	Yes for E1, E3–E5; E2 small MILP

7. Solution Methodology

7.1 Techniques in the reviewed studies

- **Article A:** an exact mixed-integer model solved using a commercial branch-and-bound solver (Gurobi), compared to an iterated tabu search algorithm. A case study using four cities for conventional vehicles and BEVs, and a sensitivity analysis over capacity, maximum travel time, service time and BEV range.
- **Article B:** a MILP characterisation of the consolidation problem, plus a custom metaheuristic designed to handle large instances where exact approaches are impractical, tested on real retailer data.

7.2 Classical methods for the transportation problem

Table 6. Classical solution steps.

Stage	Methods	Idea
Initial feasible solution	North-West Corner (NWC)	Fill cells from the top-left, ignoring cost
	Least Cost Method (LCM)	Fill the cheapest remaining cell first
	Vogel's Approximation (VAM) (Reinfeld & Vogel, 1958)	Fill the row or column with the largest regret (difference between its two cheapest cells)

Stage	Methods	Idea
Optimality test and improvement	MODI / u–v method; stepping-stone (Charnes & Cooper, 1954)	Compute reduced costs $c_{ij} - u_i - v_j$; if any is negative, move flow around a closed loop
Direct solution	Simplex-based LP solvers (Dantzig, 1951)	Solve the LP; duals give u and v

7.3 Methods used for the extensions

The models E1, E3, E4 and E5 are linear programmes and are solved exactly by the HiGHS solver (Huangfu & Hall, 2018) via SciPy. E2 involves integer van counts and is solved exactly as a small mixed-integer programme with the HiGHS branch-and-bound. E3 is solved in its extensive form and its value is measured with two standard indicators: the value of the stochastic solution (VSS) that compares the stochastic plan with a plan designed for average demand, and the expected value of perfect information (EVPI) that compares it with the cost with the demand perfectly known up front (Birge & Louveaux, 2011). E4 uses the ε-constraint method with nine evenly spaced emission limits.

7.4 Validation and software

Exact models are validated by optimality conditions (all reduced costs non-negative), by recomputing the objective values by hand, and by verifying that each extension reproduces the base case when its new feature is switched off (e.g., E2 with $w = 0$ equals FIFO, E5 with $r_j = 0$ equals the base TP). Heuristics are validated against exact solutions on small instances, as done in Article A. All computations are carried out with Python with NumPy and SciPy; other common solvers are CPLEX, Gurobi, LINGO and Excel Solver.

8. Analysis I: Solving the Base Fulfilment-to-Hub Model

All data in Sections 8 and 9 are hypothetical and chosen for teaching. They are not drawn from any company. Parameters were set to plausible orders of magnitude for an Indian metropolitan network so that the results are interpretable, and every number can be reproduced from Appendix A.

8.1 Instance

Three fulfilment centres (FC1 to FC3) supply four city hubs (H1 to H4). Costs are INR per parcel; supply and demand are parcels per day.

Table 7. Cost matrix, supply and demand.

	H1	H2	H3	H4	Supply
FC1	14	22	31	27	1,200
FC2	25	12	20	33	900
FC3	30	24	15	18	700
Demand	800	700	600	500	2,800 / 2,600

Total supply is 2,800 and total demand is 2,600, so a dummy hub with demand 200 and zero costs balances the problem. A balanced 3 × 5 problem has $m + n - 1 = 7$ basic cells in a non-degenerate basic solution.

8.2 Initial feasible solutions

North-West Corner. Starting at FC1–H1 and moving right or down gives the plan in Table 8. It uses only six cells, one fewer than $m + n - 1$, so the solution is degenerate; a zero allocation (ε) must be placed in an independent cell before the MODI test can be run. Its cost is INR 44,600 per day.

Table 8. North-West Corner plan (parcels per day).

	H1	H2	H3	H4	Dummy
FC1	800	400	–	–	–
FC2	–	300	600	–	–
FC3	–	–	–	500	200

Cost = 800×14 + 400×22 + 300×12 + 600×20 + 500×18 = INR 44,600.

Least Cost Method. Taking the cheapest real lanes first (FC2–H2 at 12, FC1–H1 at 14, FC3–H3 at 15, FC3–H4 at 18, then FC1–H4 at 27) gives a cost of INR 41,200. It is also degenerate (six cells) and leaves 200 parcels of FC2 capacity idle while FC1 ships 400 parcels on its expensive H4 lane.

Table 9. Least Cost plan (parcels per day).

	H1	H2	H3	H4	Dummy
FC1	800	–	–	400	–
FC2	–	700	–	–	200
FC3	–	–	600	100	–

Vogel's Approximation. VAM computes, for each row and column, the difference between its two cheapest cells and allocates in the line with the largest difference. Table 10 traces every step. VAM produces a non-degenerate solution with seven cells and a cost of INR 42,200.

Table 10. Vogel's Approximation, step by step.

Step	Row penalties	Column penalties	Largest	Cell filled	Qty
1	FC1 14, FC2 12, FC3 15	H1 11, H2 10, H3 5, H4 9, Dummy 0	FC3	FC3-Dummy	200
2	FC1 8, FC2 8, FC3 3	H1 11, H2 10, H3 5, H4 9	H1	FC1-H1	800
3	FC1 5, FC2 8, FC3 3	H2 10, H3 5, H4 9	H2	FC2-H2	700
4	FC1 4, FC2 13, FC3 3	H3 5, H4 9	FC2	FC2-H3	200
5	FC1 4, FC3 3	H3 16, H4 9	H3	FC3-H3	400
6	FC1 27, FC3 18	H4 9	FC1	FC1-H4	400
7	FC3 18	H4 18	FC3	FC3-H4	100

One of the surprising results is that the VAM, which typically outperforms other heuristics as a starting solution, has been improved by a LCM application. This is because the dummy column required by the degeneracy creates a row penalty of 15 for FC3 in step 1, which causes VAM to allocate the first available supply from the cheapest source (FC3) to the dummy rather than to H4, which is actually the cheapest option. Various textbooks suggest eliminating the dummy column when calculating penalties for the VAM, which would resolve the issue. In our case, the solution would be the same, which indicates that the error was not caused by the dummy but inherent to the VAM's greedy approach. Hence, we conclude that any start solution heuristics should be followed by an optimality test.

8.3 Optimality test with the MODI method

Applying MODI to the VAM plan, set $u_1 = 0$ and solve $u_i + v_j = c_{ij}$ on the seven used cells. This gives $u = (0, -4, -9)$ and $v = (14, 16, 24, 27, 9)$. The reduced cost of FC1–Dummy is $0 - 0 - 9 = -9$ and that of FC2–Dummy is $0 + 4 - 9 =$

−5, so the plan is not optimal. FC1–Dummy enters the basis. The closed loop is FC1–Dummy (+), FC3–Dummy (−), FC3–H4 (+), FC1–H4 (−), and the largest feasible shift is θ = min(200, 400) = 200 parcels. After the shift, FC3 sends 300 parcels to H4 instead of 100, FC1's H4 shipment falls from 400 to 200, and FC1, not FC3, keeps the 200 parcels of spare capacity. Cost falls by 9 × 200 = INR 1,800, to INR 40,400.

Repeating the test on the new plan gives u = (0, −4, −9) and v = (14, 16, 24, 27, 0). Table 11 shows that every reduced cost on an unused cell is strictly positive, so the plan is optimal and also unique. The LP solver returns the same plan and cost.

Table 11. Optimal plan with MODI reduced costs (used cells show parcels in brackets; other cells show the reduced cost $c - u - v$).

	H1 (v = 14)	**H2 (v = 16)**	**H3 (v = 24)**	**H4 (v = 27)**	**Dummy (v = 0)**
FC1 (u = 0)	[800]	+6	+7	[200]	[200]
FC2 (u = −4)	+15	[700]	[200]	+10	+4
FC3 (u = −9)	+25	+17	[400]	[300]	+9

Optimal cost: 800×14 + 200×27 + 700×12 + 200×20 + 400×15 + 300×18 = **INR 40,400 per day**. Compared with the starting rules, the optimum is INR 4,200 (9.4%) cheaper than NWC, INR 800 (1.9%) cheaper than LCM and INR 1,800 (4.3%) cheaper than VAM.

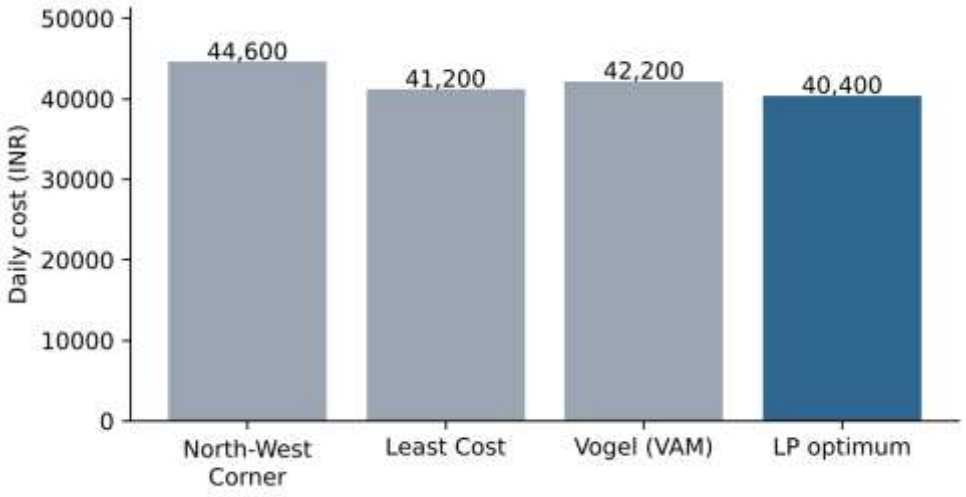


Figure 1. Daily cost of the three starting rules and the optimum (INR).

8.4 Shadow prices and cost ranging

The MODI numbers have an economic interpretation for the managers. With FC1 as the reference (its capacity is spare, so its capacity has zero value at the margin), the duals tell us:

- **Value of FC capacity.** An extra parcel of daily capacity at FC2 is worth INR 4, and at FC3 INR 9. Re-solving with an extra 100 parcels of capacity confirms the savings of INR 400 and INR 900. An expansion of FC3 is worth more than twice an expansion of FC2 because it frees up the costly H4 lane of FC1.
- **Marginal cost of demand.** An extra parcel of demand costs INR 14, 16, 24 and 27 at the four hubs. These are the true marginal costs of pricing delivery or setting minimum-order thresholds by area: different from the cheapest lane cost into each hub (for H3, the cheapest lane costs 15, but an extra H3 parcel costs 24 once the knock-on effects on other hubs are counted).
- **Cost ranging.** If only the FC1–H4 rate changes by δ, the reduced costs on the unused cells become 6 − δ, 7 − δ, 15 + δ, 10, 4 + δ, 25 + δ, 17 and 9 + δ. All these are still non-negative for −4 ≤ δ ≤ 6, so the plan is still optimal while

that rate lies between INR 23 and INR 33 per parcel: a carrier quoting a price above INR 33 on that lane will cause a re-plan; those quoting below INR 23 would win more volume.

8.5 Scenario analysis

Table 12. Scenarios on the base model.

Scenario	Daily cost (INR)	Change	Explained by
Base case	40,400	–	–
All lane costs +20% (fuel shock)	48,480	+20.0%	Plan unchanged; cost scales
FC2 capacity falls from 900 to 700	41,200	+2.0%	Shadow price 4 × 200
H4 demand rises from 500 to 700 (festival surge)	45,800	+13.4%	Marginal cost 27 × 200

The shadow prices predict each change exactly. The surge case is the most instructive because total demand reaches 2,800 and utilizes all spare capacity, which a 7.7% increase in volume raises cost by 13.4%. This is the base case's first sign of the uncertainty gap (G3), which Section 9.3 analyzes properly.

9. Analysis II: Extending the Transportation Model to Fill the Gaps

9.1 Shared depots between carriers (G1)

Set-up. Two carriers, P1 and P2, each operate two depots in a hypothetical Mumbai-like city and serve their own customers in six delivery zones. Distances are Manhattan distances (the same metric Article A uses) on a city grid, and the line-haul cost per parcel is INR 6 of handling plus INR 0.90 per km. Depot capacities are 1,700, 1,500, 1,700 and 1,300 parcels per day. Table 13 gives distances and demand.

Table 13. Depot-to-zone distances (km) and daily demand by carrier (parcels).

Depot	Andheri	Bandra	Dadar	Powai	Thane	Navi Mumbai
P1-North (cap 1,700)	12	23	27	14	14	34
P1-South (cap 1,500)	20	11	5	26	42	26
P2-East (cap 1,700)	28	27	31	14	14	10
P2-West (cap 1,300)	8	5	13	14	30	34
Demand P1	520	610	450	380	300	240
Demand P2	340	280	390	460	520	560

Results. Solving each carrier's TP separately would cost INR 42,864 for P1 and INR 40,959 for P2, or INR 83,823 per day in total. By contrast, solving one pooled TP would cost INR 74,013. This represents a saving of INR 9,810 per day (or 11.7%). The number of parcel-kilometres would fall from 59,470 to 48,570 (a reduction of 18.3%), and the average line-haul distance per parcel would fall from 11.8 km to 9.6 km.

Table 14. Pooled allocation (parcels per day, both carriers combined).

Depot	Andheri	Bandra	Dadar	Powai	Thane	Navi Mumbai	Used
P1-North	450	–	–	840	–	–	1,290
P1-South	–	–	840	–	–	–	840
P2-East	–	–	–	–	820	800	1,620
P2-West	410	890	–	–	–	–	1,300

Interpretation. The saving is a pure allocation effect: no route has changed, only which depot feeds which zone. In the siloed plans each carrier must cover the whole city from its own two depots, so P1 serves Navi Mumbai from its South depot 26 km away and P2 serves Dadar from its West depot 13 km away. In the pooled plan each zone is fed from the nearest depot with spare capacity, whoever owns it: P1-South takes all of Dadar (5 km), P2-West takes Bandra (5 km) and fills its capacity with part of Andheri, P1-North takes the rest of Andheri and Powai, and P2-East takes Thane and Navi Mumbai (14 km and 10 km). This splits the two effects that Article A reports together. In our instance about 18% of distance is saved before any routing is improved, which suggests that a substantial share of the 38.8% VMT saving in Austin could come from the allocation layer alone, with the remainder coming from denser routes.

Capacity sensitivity. When every depot's capacity is reduced by the same proportion the pooled saving shrinks: 11.7% at full capacity, 11.4% at 95%, 11.1% at 90% and 10.6% at 86%. Pooling pays most when there is spare capacity to redirect, so its value is highest outside peak periods.

Sharing the gain. Collaboration is only stable if each carrier is better off than alone. A simple proportional rule allocates the pooled cost in proportion to stand-alone costs: P1 would pay INR 37,848 (saving 5,016) and P2 INR 36,165 (saving 4,794). Both gain, so the coalition is individually rational. With more carriers rules from cooperative game theory such as the Shapley value are preferred (Guajardo & Rönnqvist, 2016). The point for practice is that the transportation model gives the numbers that any such agreement needs, quickly and transparently.

9.2 Consolidation with a price on waiting (G2)

Set-up. One delivery zone is served by the order stream as shown in Table 15 over a period of 14 days (1,526 parcels). Vans carry $Q = 100$ parcels, each van trip costs F = INR 1,800 and each parcel costs c = INR 6 to handle. Under FIFO every day's orders are left out the same day. Model E2 allows for an order to be left out up to two days after its release, with a penalty of h per parcel-day representing the incentive paid to the customer or the goodwill lost.

Table 15. Daily orders and van loads under FIFO and a one-day window (h = INR 4).

Day	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Orders	138	74	162	95	57	211	46	129	83	118	67	190	52	104
Vans, FIFO	2	1	2	1	1	3	1	2	1	2	1	2	1	2
Load, 1-day	100	100	100	100	100	200	83	100	100	100	97	190	0	156
Vans, 1-day	1	1	1	1	1	2	1	1	1	1	1	2	0	2

Results. FIFO requires 22 van trips at 69.4% average utilization and costs INR 48,756. By allowing 1 day of waiting with h = INR 4, van trips are reduced to 16 and utilization jumps to 95.4%. Costs drop to INR 37,956 (22.2%), with a penalty of INR 1,516 for waiting, resulting in a total of INR 39,472, a saving of 19.0%. Average waiting time for a parcel is only 0.25 days, since the model only holds back the parcels needed to avoid a half-van and ships the rest the same day.

Table 16. Net saving vs. FIFO by delay penalty and window length (van trips in brackets; FIFO = 22 trips).

Delay penalty *h* per parcel-day	1-day window	2-day window	3-day window
INR 0	22.2% (16)	22.2% (16)	22.2% (16)
INR 4	19.0% (16)	19.1% (16)	19.1% (16)
INR 10	14.7% (17)	14.7% (17)	14.7% (17)
INR 20	10.9% (17)	10.9% (17)	10.9% (17)
INR 40	5.0% (20)	5.0% (20)	5.0% (20)
INR 60	3.8% (20)	3.8% (20)	3.8% (20)

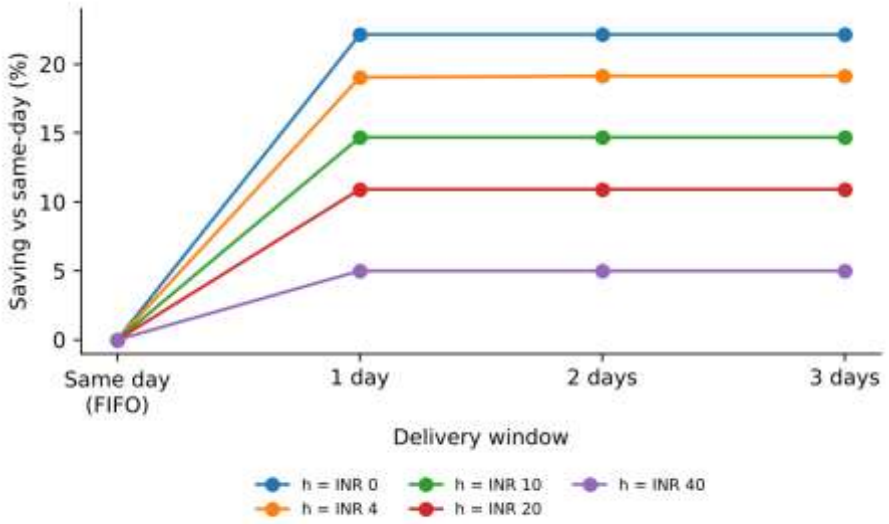


Figure 2. Net saving from consolidation by window length, for five levels of delay penalty.

Interpretation. Three results fill the gap left by Article B; first, consolidation remains profitable even with a high penalty for waiting: at INR 40 per parcel-day, more than six times the handling cost, it still saves 5.0%, since the benefit is lumpy (a whole van at INR 1,800) while the penalty only applies to the few parcels held back. Second, the penalty changes how much consolidation is done: as *h* increases the optimal van-trips moves from 16 to 17 to 20. Third, in this instance almost all the gain comes from the first day; a second or third add at most 0.1 percentage points. A retailer could therefore offer a one-day "eco-delivery" option and capture nearly all the benefits. Article B's larger savings (up to 52%) stem from consolidating across many zones and routes at once, in which longer windows also generate denser routes; our single-zone model isolates the pure vehicle-filling effect and shows that it alone is worth roughly one-fifth of cost.

9.3 Planning for uncertain and festive demand (G3)

Set-up. Using the base network, the planner needs to book FC capacity (staff and shifts) one day ahead at INR 4.0, 5.0 and 4.5 per parcel for FC1-FC3, up to a physical maximum of 1,500, 1,100 and 900 parcels. Demand then materializes as one of three scenarios, scaled from the base demand of 2,600 parcels. Parcels in excess of the booked capacity are sent to a 3PL carrier at an increased cost of INR 70 per parcel.

Table 17. Demand scenarios.

Scenario	Probability	Demand multiplier	Total parcels
Normal day	0.6	1.00	2,600
Sale event	0.3	1.25	3,250

Scenario	Probability	Demand multiplier	Total parcels
Festive peak	0.1	1.60	4,160
Expected		**1.135**	**2,951**

Table 18. Booking decisions and expected daily cost of three planning approaches (INR).

Approach	Booked capacity FC1 / FC2 / FC3	Normal	Sale	Festive	Expected
Plan for a normal day	800 / 900 / 900	50,350	93,475	156,450	73,898
Plan for average demand	951 / 1,100 / 900	51,954	77,223	137,998	68,139
Stochastic plan	1,500 / 1,100 / 900	54,150	65,925	112,210	63,488
Perfect foresight (wait-and-see)	Chosen after demand is known	–	–	–	60,908

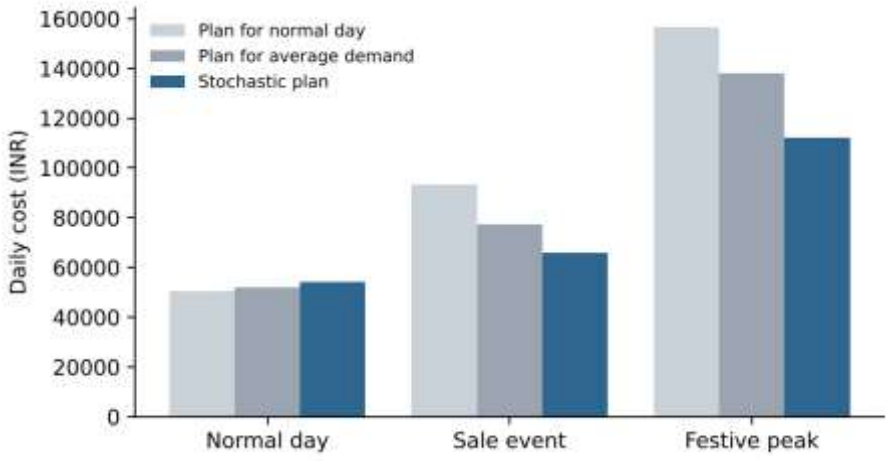


Figure 3. Daily cost in each scenario under the three planning approaches.

Results. The stochastic plan books all available capacity. It costs INR 3,800 more than the normal-day plan on a normal day, but far less on sale and festive days, and its expected cost is INR 63,488. The value of the stochastic solution is INR 4,651 per day against planning for average demand (6.8% of that plan's cost) and INR 10,409 per day against planning for a normal day. The expected value of perfect information is INR 2,580 per day, which is the most a firm should pay for a better demand forecast in this setting.

Interpretation. Planning for the average sounds prudent but is expensive, because the cost of being short (INR 70 per parcel) is many times the cost of booking spare capacity (INR 4–5). The stochastic model makes this asymmetry visible and prices it. Even the stochastic plan leaves 660 parcels to the 3PL on festive days because physical capacity is only 3,500 parcels; the model therefore also signals when a permanent capacity expansion, rather than better booking, is needed. Neither reviewed study can answer this question because both plan for one known set of orders.

9.4 Cost versus emissions (G4)

Set-up. Table 19 presents illustrative distances for each lane in the base network. Diesel vans emit an assumed 0.030 kg CO_2 per parcel-km while electric or CNG vans emit 0.012 kg, but the cleaner option incurs costs 18% higher than the cheapest one per parcel. The ε-constraint approach explores the frontier between the cost minimal solution and the emission minimal solution.

Table 19. Illustrative lane distances (km).

	H1	H2	H3	H4
FC1	18	46	72	60
FC2	52	15	38	80
FC3	70	50	22	30

Table 20. Points on the cost–emissions frontier.

Emissions (kg CO_2/day)	Cut vs. min-cost	Cost (INR/day)	Cost increase	Parcels on clean vans	Marginal cost (INR/kg)
1,869.0	0%	40,400	0.0%	0	–
1,728.8	7%	41,031	1.6%	130	4.5
1,588.7	15%	41,711	3.2%	294	4.9
1,448.5	22%	42,498	5.2%	525	5.6
1,308.3	30%	43,377	7.4%	816	6.3
1,168.1	37%	44,359	9.8%	1,185	7.0
1,028.0	45%	45,449	12.5%	1,618	7.8
887.8	52%	46,551	15.2%	2,081	7.9
747.6	60%	47,672	18.0%	2,600	8.0

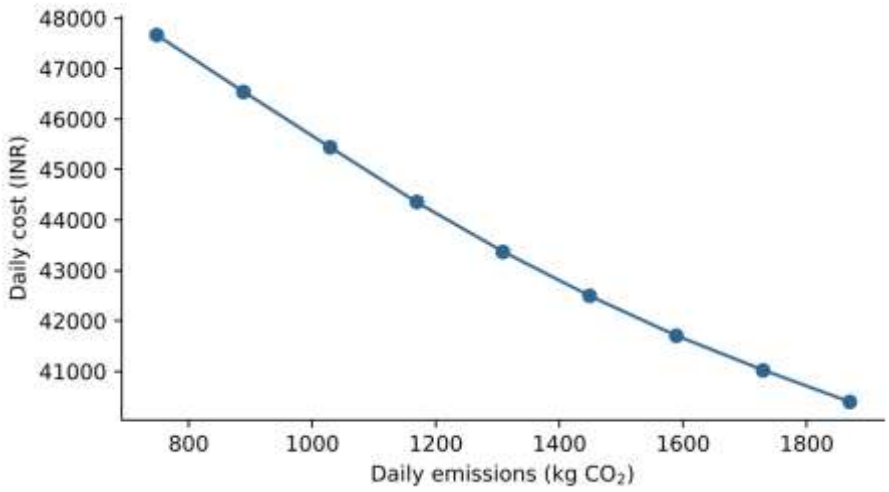


Figure 4. Pareto frontier of daily cost against daily emissions.

Results. The cost-minimal plan emits 1,869 kg CO_2 per day using no clean vans. Emissions can be reduced by 30% for a 7.4% cost increase and by 60% for an 18.0% cost increase. The marginal abatement cost increases from INR 4.5 to INR 8.0 per kg as cleaner vans are introduced to lanes where they save less carbon per rupee.

Interpretation. The frontier reveals to a manager what each step towards a sustainability target costs and in what order lanes should be converted: those offering the most kilometres per rupee of cost first (FC1-H4 is converted first in the current plan) and the shortest, such as FC1-H1 and FC2-H2, last. Across all twelve lanes, the break-even carbon price at which a lane converts to clean vans on cost considerations alone ranges from INR 4.1 to INR 8.0 per kg (INR 4,100–8,000 per tonne). Article A showed that BEVs are technically feasible on urban routes; this analysis adds the economic trade-off that determines when they are adopted.

9.5 Failed deliveries and returns to origin (G5)

Set-up. Hubs H1-H3 serve densely populated urban areas with low failure rates of 8%, 10% and 12%, while H4 serves a catchment where cash on delivery is the norm and 30% of parcels are refused and returned. These are assumptions for illustration. A failed parcel traverses back at 80% of the cost of the forward delivery ($\rho = 0.8$) and is re-sent until it is delivered. When FC capacity is inadequate, a 3PL carrier supplies parcels at INR 70 each.

Table 21. Parcels that must be shipped per day and the RTO-adjusted allocation.

Source	H1 (8%)	H2 (10%)	H3 (12%)	H4 (30%)	Total
FC1	870	–	–	330	1,200
FC2	–	778	122	–	900
FC3	–	–	316	384	700
3PL	–	–	243	–	243
Shipments required	870	778	682	714	3,043
Delivered demand	800	700	600	500	2,600

Results. To deliver 2600 parcels the network must ship 3043, i.e. more than the 2800 parcels of capacity that the FCs provide. A plan which ignores RTO, the base plan which costs INR 40,400, is therefore not just optimistic, but infeasible. The adjusted model sends about 243 parcels to the 3PL and costs INR 69,217 per day, raising the cost per delivered parcel from INR 15.54 to INR 26.62. The allocation also changes, with FC1 sending more parcels to H4 and FC3 carrying more from H3 to H4 (being the effective lane where the cost per parcel increases most).

A policy test. If prepaid incentives or address verification reduce the failure rate to 15%, the number of parcels sent falls to 2,917 and outsourcing falls to about 117 parcels and cost per day falls to INR 56,939, a saving of INR 12,279 per day (17.7%). This would provide a ceiling on the amount the retailer might spend per day on reducing failures in this catchment. Because the mechanism works through the capacity, as well as cost, it would not be visible in a routing model which takes as given the parcels to be delivered.

10. Critical Analysis and Discussion

10.1 Key findings of the reviewed studies

- **Article A:** service time followed by vehicle capacity determines the performance of the system; BEV range was rarely limiting due to the average drive distance of 72 miles compared to the assumed 80-mile range; and operating depots jointly cut VMT in Austin by 38.8%.
- **Article B:** consolidation leads to higher vehicle utilisation while reducing the distance, time and number of fleets while delivering cost savings of up to 52% for one real-world retailer.

10.2 What the gap analysis adds

Table 22. Summary of the gap-filling analysis.

Gap	Model	Main finding	Managerial implication
G1 Sharing	E1 pooled TP	11.7% lower cost, 18.3% fewer parcel-km from allocation alone	Depot sharing pays even before routes are merged; savings shrink as capacity tightens
G2 Waiting	E2 time-indexed TP	19.0% net saving with a 1-day window; extra days add almost nothing	Offer a one-day eco-delivery option; price the delay explicitly

Gap	Model	Main finding	Managerial implication
G3 Uncertainty	E3 stochastic TP	VSS INR 4,651/day; EVPI INR 2,580/day	Book capacity against scenarios, not the average
G4 Emissions	E4 bi-objective TP	30% CO_2 cut for 7.4% more cost	Convert long lanes first; frontier sets the price of targets
G5 Returns	E5 RTO-adjusted TP	Cost per delivered parcel +71%; base network infeasible	Plan capacity on shipments, not orders; failure reduction has a measurable value

10.3 Impact and implications

Both reviewed studies indicate that large savings can be realized by changing the grouping and sharing of orders rather than by simply accelerating their movement. Our analysis similarly underscores the importance of the allocation layer and provides three additional insights. First, a substantial component of the value of collaboration and consolidation can be captured in a transportation model solvable within milliseconds, enabling companies to screen consolidation-type policies before embarking on a full routing study. Second, the cost of waiting, risk of peaks and cost of returns transform the optimization metric and must be incorporated at the core of the model, rather than as an afterthought. Finally, the duals of the transportation model are directly informative: they represent the value of FC capacity, establish marginal delivery costs by area and identify which carrier quotes would alter the optimal plan.

A practical planning pipeline emerges: (1) solve a stochastic, RTO-adjusted transportation model to fix FC capacity and allocates hub volumes; (2) apply a time-indexed consolidation model at each hub to identify what leaves today; (3) route the resulting loads with a CVRP or a metaheuristic, as in Article A. Outputs of each layer feed into the next and each can be solved with the methods outlined in Section 7.

10.4 Strengths

- Exact models provide optimality guarantees and clean sensitivity analysis, and the MODI prices have business interpretation value.
- The transportation model is fast and transparent, and because it is totally unimodular any integer data will yield integer plans.
- Each extension addresses a particular gap, isolating it to examine its impact.
- Metaheuristics, as applied in both of the discussed articles, are helpful in addressing routing and consolidation to scale to realistic instances.

10.5 Limitations and assumptions

- All numbers in Sections 8 and 9 are derived from hypothetical, small examples, which illustrate the direction and magnitude of effects, but are not meant to represent realistic figures for any actual company.
- Costs scale linearly with parcel volumes, except for van fixed costs in E2, where real-life line hauls have truck-load economies and lane minimum costs.
- The consolidation model is for a single zone and does not account for gains through route-density increases, thus under-estimates the benefits that Article B reports.
- The Stochastic model is based on three discrete scenarios; real-life demand distributions would require more scenarios or a robust optimization approach.
- Emission factors, the clean-van premium, and RTO rates are illustrative numbers and should be replaced with measurements.
- The results for the U.S. cities and a Mexican retailer may not generalize to dense Indian cities with informal addressing and two-wheeler fleets.

10.6 Directions for further research

- Combine E1–E5 into one model and test which gaps interact (for instance, whether pooling reduces the value of the stochastic solution).
- Calibrate the models using real Indian data on demand peaks, failure rates by PIN code and lane costs.
- Add route-density effects to E2 by linking it to a CVRP at each hub.
- Replace the proportional sharing rule in E1 with the Shapley value or nucleolus for three or more carriers.
- Use rolling-horizon re-optimisation with live order streams, as in dynamic dispatch research (Klapp, Erera & Toriello, 2018).
- Model two-echelon networks with micro-hubs, lockers and two-wheeler or electric fleets.

11. Conclusion

This paper reviewed two journal-level applications of OR to e-commerce logistics and found that the classical transportation problem was embedded in both as the flow of packages in the delivery network (in Article A) or the temporal arrangement of consolidations (in Article B). The base fulfilment-to-hub model was solved with all three initial assignment rules, tested with the MODI method, and interpreted through shadow prices; the optimal solution was proven to be 9.4% better than with the North-West Corner method, and the VAM was not always the best starting point.

Five gaps were found in the reviewed articles: allocation benefits from shared depots, customer waiting costs, demand-side uncertainty, cost–emissions trade-offs, and returns to the origin. Five extended transportation models were found to fill these gaps: pooling depots reduced costs by 11.7% over allocation alone; a one-day delivery window reduced delivery costs by 19% net of penalties; scenario-based capacity booking reduced costs by INR 4,651 per day on average over plan-driven booking; clean vans reduced emissions by 30% for a 7.4% increase in costs; RTOs increased the cost per delivered package by 71% and rendered the base network configuration invalid. Thus, the results suggest a three-stage logistics planning process: allocate deliveries according to a transportation model with uncertain returns, consolidate the deliveries in a time-indexed model, and route them. Further research should focus on finding exact coefficients for the reviewed models and linking them into a cohesive system.

References

Birge, J. R., & Louveaux, F. (2011). *Introduction to stochastic programming* (2nd ed.). Springer.

Boysen, N., Fedtke, S., & Schwerdfeger, S. (2021). Last-mile delivery concepts: A survey from an operational research perspective. *OR Spectrum, 43*(1), 1–58.

Charnes, A., & Cooper, W. W. (1954). The stepping stone method of explaining linear programming calculations in transportation problems. *Management Science, 1*(1), 49–69.

Cleophas, C., Cottrill, C., Ehmke, J. F., & Tierney, K. (2019). Collaborative urban transportation: Recent advances in theory and practice. *European Journal of Operational Research, 273*(3), 801–816.

Cokyasar, T., Subramanyam, A., Larson, J., Stinson, M., & Sahin, O. (2023). Time-constrained capacitated vehicle routing problem in urban e-commerce delivery. *Transportation Research Record, 2677*(2), 190–203. https://doi.org/10.1177/03611981221124592

Dantzig, G. B. (1951). Application of the simplex method to a transportation problem. In T. C. Koopmans (Ed.), *Activity analysis of production and allocation* (pp. 359–373). Wiley.

Dantzig, G. B., & Ramser, J. H. (1959). The truck dispatching problem. *Management Science, 6*(1), 80–91.

Guajardo, M., & Rönnqvist, M. (2016). A review on cost allocation methods in collaborative transportation. *International Transactions in Operational Research, 23*(3), 371–392.

Haimes, Y. Y., Lasdon, L. S., & Wismer, D. A. (1971). On a bicriterion formulation of the problems of integrated system identification and system optimization. *IEEE Transactions on Systems, Man, and Cybernetics, SMC-1*(3), 296–297.

Hitchcock, F. L. (1941). The distribution of a product from several sources to numerous localities. *Journal of Mathematics and Physics, 20*(1–4), 224–230.

Huangfu, Q., & Hall, J. A. J. (2018). Parallelizing the dual revised simplex method. *Mathematical Programming Computation, 10*(1), 119–142.

Klapp, M. A., Erera, A. L., & Toriello, A. (2018). The dynamic dispatch waves problem for same-day delivery. *European Journal of Operational Research, 271*(2), 519–534.

Koopmans, T. C. (1949). Optimum utilization of the transportation system. *Econometrica, 17*(Supplement), 136–146.

Muñoz-Villamizar, A., Velázquez-Martínez, J. C., & Caballero-Caballero, S. (2024). A large-scale last-mile consolidation model for e-commerce home delivery. *Expert Systems with Applications, 235*, 121200. https://doi.org/10.1016/j.eswa.2023.121200

Reinfeld, N. V., & Vogel, W. R. (1958). *Mathematical programming*. Prentice-Hall.

Savelsbergh, M., & Van Woensel, T. (2016). 50th anniversary invited article—City logistics: Challenges and opportunities. *Transportation Science, 50*(2), 579–590.

Taha, H. A. (2017). *Operations research: An introduction* (10th ed.). Pearson.

Toth, P., & Vigo, D. (Eds.). (2014). *Vehicle routing: Problems, methods, and applications* (2nd ed.). SIAM.

Appendix A. Reproducibility and parameter summary

All models were built in Python 3 with NumPy and SciPy and solved with HiGHS through scipy.optimize.linprog (linear models) and scipy.optimize.milp (model E2). The base model has 15 flow variables after balancing, three supply rows and five demand rows; the LP solution and its equality-constraint marginals give the plan and the MODI prices reported in Section 8. Each extension is a separate function that changes one element of the base model, so switching the new feature off reproduces the base result.

Table 23. Parameter summary for the extended models.

Model	Key parameters (all hypothetical)
E1 Pooled TP	2 carriers, 4 depots, 6 zones; cost = INR 6 + 0.90 × km; depot capacities 1,700 / 1,500 / 1,700 / 1,300
E2 Consolidation	14 days, 1,526 parcels; Q = 100; F = INR 1,800 per van; c = INR 6; h from 0 to 60; w from 0 to 3 days
E3 Stochastic TP	Booking INR 4.0 / 5.0 / 4.5; maxima 1,500 / 1,100 / 900; 3PL INR 70; scenarios 0.6 / 0.3 / 0.1 at ×1.00 / 1.25 / 1.60
E4 Bi-objective TP	Diesel 0.030 and clean 0.012 kg CO_2 per parcel-km; clean premium 18%; nine ε levels
E5 RTO-adjusted TP	RTO 8% / 10% / 12% / 30%; return leg 80% of forward cost; 3PL INR 70