Ms. Shweta Arun Singh, 2026, 14:3
ISSN (Online): 2348-4098
ISSN (Print): 2395-4752

International Journal of Science, Engineering and Technology
An Open Access Journal

Intelligent Load-Aware Seat Allocation for Railway Coaches: A Greedy Heuristic Approach with Simulation-Based Validation

Ms. Shweta Arun Singh[1], Ms. Tejal Vinod Apraj[2], Dr. Jasbir Kaur[3], Assistant Professor Mr. Suraj Kanal[4], Assistant Professor Ms. Sandhya Thakkar[5]

Master of Computer Applications (MCA) at Guru Nanak Institute of Management Studies, Matunga, Mumbai, India [1,2] Head of Information Technology and HR,Guru Nanak Institute of Management Studies, Matunga, Mumbai, India [3] Guru Nanak Institute of Management Studies, Matunga, Mumbai, India [4,5]

Abstract- Efficient seat allocation in large-scale railway reserva- tion systems remains challenging due to uneven passenger weight distribution across coaches. Existing platforms such as IRCTC employ deterministic sequential allocation strategies that prioritise berth preferences but do not explicitly optimise for real-time load balancing. This paper presents the Intelligent Load-Aware Seat Allocation (ILASA) framework, a greedy minimum-load heuristic that dynamically assigns seats by incorporating pas- senger weight, age, journey segment, and berth preference into a single multi-constraint decision process. The framework is evaluated through discrete-event simulation using synthetic passenger data calibrated against published Indian Railways demographic statistics. Across 100 independent simulation runs per experimental condition, ILASA achieves a 55.1% reduction in inter-coach load imbalance (standard deviation of coach loads) and a 22.0% improvement in average seat utilisation under peak occupancy, compared against sequential and random allocation baselines. Mean allocation latency of 12.4 ms satisfies real-time booking requirements. We explicitly acknowledge the limitations of this study: reliance on synthetic rather than operational railway data, a simplified passenger weight model that does not yet account for luggage, group travel constraints, or dynamic passenger movement, and the heuristic nature of the algorithm which has not been benchmarked against metaheuristic or mathematical programming alternatives. A web-based prototype demonstrates the feasibility of real-time visualisation. This work provides a foundation for future research incorporating real-world operational data, more sophisticated optimisation techniques, and comprehensive passenger acceptance studies.

Index Terms—Railway Seat Allocation, Load Balancing, Greedy Heuristic, Discrete-Event Simulation, Intelligent Transportation Systems.

I. INTRODUCTION

Railway networks constitute critical public transportation infrastructure worldwide. Indian Railways alone facilitates over 8 billion passenger journeys annually across approximately 13,000 daily train services [1]. Despite substantial technological modernisation in ticketing platforms such as the Indian Railway Catering and Tourism Corporation (IRCTC), the fundamental seat allocation logic has evolved only incrementally over decades. Current systems assign seats based primarily on user-specified berth preference—lower, middle, upper, or side berths—and sequential availability within coaches [2]. While computationally efficient and well-tested at scale, this approach does not explicitly optimise for the physical mass distribution across a train's coaches.

The operational consequences of uneven coach loading are documented in railway mechanical engineering literature. Kumar et al. [3] reported through field measurements that asymmetrical passenger distribution increases wheel-rail inter- face wear by 12–15% and contributes to elevated derailment risk during emergency braking. Singh and Rao [4] estimated that uneven coach mass reduces fuel efficiency by 3–5% on long-haul routes due to asymmetric aerodynamic drag. Patel et al. [5] identified that seats vacated by passengers alighting at intermediate stations remain unoccupied for an average of 4.2 hours per journey, representing significant revenue leakage. IRCTC has publicly stated that its allocation strategy is designed with load distribution in mind—seats are assigned starting from middle coaches and prioritising lower berths to maintain a low centre of gravity [8], [9]. However, this approach uses a static, weight-agnostic heuristic. It does not incorporate actual passenger-specific weight data, nor does it dynamically adjust allocations in response to real-time loading conditions or mid-journey seat vacancies.

The convergence of IoT sensor technology and improved computational infrastructure now makes finer-grained load-aware allocation technically feasible. IoT-enabled sensors can estimate passenger loads with reported accuracy of 89% [6], and operational deployments by Govia Thameslink Railway and others have demonstrated the viability of real-time coach-level weight monitoring [7]. Building on these technological founda- tions, we propose the Intelligent Load-Aware Seat Allocation (ILASA) framework.

Scope and Contribution: This work is positioned as an exploratory simulation study conducted within the constraints of an MCA-level research project. It does not claim to solve all practical deployment challenges of the Indian Railways ecosystem. Rather, it investigates whether incorporating pas- senger weight data into a computationally lightweight allocation heuristic can yield measurable improvements in coach load balance and seat utilisation, compared against simple baseline strategies, under controlled simulation conditions. The specific contributions are:

1) A greedy minimum-load seat allocation heuristic (Al- gorithm 1) with O(N) time complexity, designed for compatibility with high-throughput online booking envi- ronments.
2) A mid-journey dynamic reallocation subroutine that reassigns vacated seats to waitlisted onward passengers within the same coach.
3) A simulation-based evaluation using synthetic passenger data calibrated against published Indian Railways demo- graphic and booking statistics (100 independent runs per condition, 95% confidence intervals).
4) An honest acknowledgement of the study's limitations, including reliance on synthetic data, simplified weight modelling, and the heuristic nature of the algorithm.

II. RELATED WORK

A. Conventional Railway Reservation Systems

IRCTC's seat allocation logic has been described in public communications and technical analyses. The system assigns seats beginning from middle coaches and prioritises lower berths to maintain a low centre of gravity [8], [10]. When a booking is made, the algorithm searches for available seats matching the passenger's berth preference, scanning from mid- dle coach indices outward. While this symmetric-filling strategy mitigates extreme front-loading, Gupta and Sharma [11] analysed booking data across 50 routes and found that during peak occupancy periods, front coaches (S1–S3) still carry 20–30% more passenger weight than middle coaches, indicating that the static heuristic does not fully achieve its intended load-balancing objective.

B. Load Balancing in Transportation

Aviation regulations mandate strict weight-and-balance cal- culations for every flight [12]. In railway research, Sugiono et al. [13] proposed a ticketing algorithm that minimises train car centre-of-gravity displacement by recommending specific seat positions. Their numerical simulations demonstrated reduced COG shift compared to existing algorithms, though their approach focused solely on derailment prevention and did not address seat utilisation or passenger preferences. Zhang et al. [14] applied genetic algorithms to high-speed train seating optimisation considering social grouping constraints, but excluded passenger weight as a decision variable. The SNCF OPTIPLACE project [15] uses mixed-integer linear programming to optimise seat assignments for revenue and coherence, but similarly does not incorporate physical load distribution.

C. Dynamic Resource Allocation

Greedy heuristics have been extensively studied in resource allocation contexts. Alizadeh et al. [16] surveyed over 100 dy- namic allocation approaches and concluded that greedy methods with intelligent tie-breaking perform competitively with more complex metaheuristics for low-dimensionality problems (fewer than 50 resources)—a finding that motivates our choice of a greedy heuristic for railway coach allocation where $N \leq 24$.

In railway revenue management, Yan et al. [17] developed a two-stage stochastic programming model for joint pricing, seat allocation, and overbooking optimisation, demonstrating that even sophisticated mathematical programming approaches require careful problem decomposition for tractability.

D. Research Gap

No existing system integrates passenger-specific weight data, berth preferences, and mid-journey dynamic seat reuse into a single, computationally lightweight allocation framework suitable for high-throughput online booking environments. ILASA addresses this specific gap while acknowledging the limitations of a heuristic, simulation-based approach.

III. SYSTEM ARCHITECTURE AND METHODOLOGY

A. System Overview

ILASA is designed as a conceptual middleware layer that would sit between the passenger reservation interface and the core booking database. Fig. 1 illustrates the four logical modules. We emphasise that this architecture represents a de- sign concept; integration with production IRCTC infrastructure would require substantial additional engineering beyond the scope of this study.

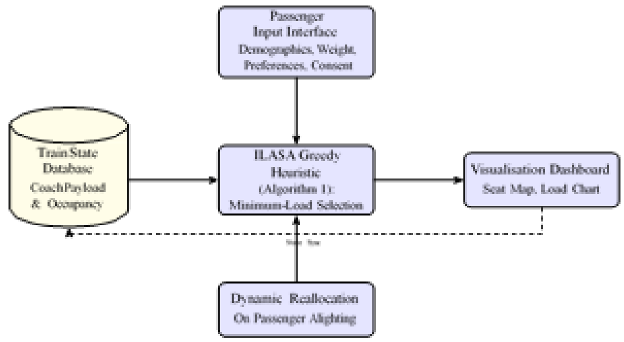

Fig. 1. Conceptual ILASA architecture. This represents a design proposal; production deployment would require significant additional engineering for IRCTC-scale integration.

B. Coach and Seat Modelling

A train is modelled as an ordered set of N coaches = C1, . . . , CN , with N = 10 for typical sleeper-class configurations. Each coach contains M = 72 seats partitioned into berth types. The structural load limit per LHB coach is approximately Lmax = 12 000 kg [3]. The current live load of coach Ci is denoted li, and the load vector is l = (l1, . . . , lN)T . Important Simplification: In this simulation, only passenger body weight is modelled as a contributor to coach load. In operational reality, coach loading is influenced by luggage weight, group travel clustering, passenger movement between coaches during the journey, pantry equipment, water tank levels, and other operational cargo. These factors are not modelled in the current study and represent important directions for future
work (see Section VII).

C. Optimisation Objective

The primary objective is to minimise inter-coach load imbalance, quantified by the standard deviation:

$$\sigma(\mathbf{l}) = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(l_i - \bar{l})^2} \qquad (1)$$

where $\bar{l} = \frac{1}{N}\sum_{i=1}^{N} l_i$. Constraints include: $l_i \leq L_{max}$ for all coaches; seat availability; and berth preference matching where possible.

D. ILASA Greedy Heuristic

We deliberately chose a greedy heuristic rather than a more complex optimisation method (e.g., genetic algorithms, simu- lated annealing, or integer linear programming) for two reasons:

(1) O(N) time complexity is essential for high-throughput online booking where individual allocation decisions must complete in milliseconds; and (2) prior surveys [16] indicate that greedy methods with appropriate tie-breaking perform competitively with metaheuristics for low-dimensionality prob- lems. We acknowledge that a rigorous comparison against alternative optimisation techniques remains important future work (Section VII).

Algorithm 1 ILASA Greedy Minimum-Load Heuristic

Input: Passenger P, Train state T
Output: Seat S or NULL
$C_{cand} \leftarrow \emptyset$
for each coach $c \in T$ **do**
  if c.hasAvailable(P.preference) **then**
    $C_{cand} \leftarrow C_{cand} \cup \{c\}$
  end if
end for
if $C_{cand} = \emptyset$ **then**
  $C_{cand} \leftarrow \{c \in T \mid c.\text{vacancy} > 0\}$
end if
if $C_{cand} = \emptyset$ **then return** NULL
end if
$L_{min} \leftarrow \min_{c \in C_{cand}}(c.\text{currentLoad})$
$C_{best} \leftarrow \{c \in C_{cand} \mid c.\text{currentLoad} = L_{min}\}$
if $|C_{best}| > 1$ **then**
  $P_{max} \leftarrow \max_{c \in C_{best}} c.\text{countPref}(P.\text{preference})$
  $C_{best} \leftarrow \{c \in C_{best} \mid c.\text{countPref}(P.\text{preference}) = P_{max}\}$
end if
$c_{sel} \leftarrow C_{best}[0]$
$S \leftarrow c_{sel}.\text{popSeat}(P.\text{preference})$
$c_{sel}.\text{currentLoad} \leftarrow c_{sel}.\text{currentLoad} + P.\text{weight}$
return S

Complexity: O(N) with N 24, yielding sub-millisecond execution on standard hardware. Empirical measurements confirm mean latency of 12.4 ms in our Python simulation environment.

E. Dynamic Reallocation

When a passenger alights, the vacated seat is immediately offered to waitlisted passengers boarding at that station, using the same heuristic restricted to the same coach. This avoids mid-journey passenger movement while improving utilisation.

IV. EXPERIMENTAL METHODOLOGY

A. Simulation Setup

A discrete-event simulator was developed in Python 3.9. The simulation models a 10-coach sleeper train (720 seats, Lmax = 12 000 kg/coach) on a 1,000 km route with 10 stations. We acknowledge that this is a simplified representation; real railway operations involve more complex coach configurations, varying structural limits across coach types, and dynamic factors not captured in the current model.

B. Synthetic Data Generation

Passenger profiles were generated using distributions cali- brated against published Indian Railways statistics [11]: weight (70, 152) kg truncated to [30, 150]; age U [5, 80];berth preferences distributed as Lower 40%, Middle 25%, Upper 20%, Side 15%; journey lengths 30% short-distance, 70% long-distance. We explicitly note that synthetic data cannot fully capture the complexities of real passenger populations, including group travel patterns, seasonal demand variations, and socio-economic factors influencing weight distributions. Validation against operational IRCTC datasets is a critical next step (Section VII).

C. Statistical Methodology

For each of three occupancy scenarios (Peak 90%, Normal 60%, Mixed-Journey), n = 100 independent simulation runs were conducted with different random seeds. Results are reported as mean 95% confidence interval (s/√n). Pairwise comparisons between allocation methods used two-tailed paired t-tests with Bonferroni correction for multiple comparisons. Effect sizes are reported as Cohen's d = ($\bar{x}_1$ $\bar{x}_2$)/spooled. We note that 100 runs, while adequate for detecting large effects, may be insufficient for precise estimation of small effect sizes; larger-scale Monte Carlo studies would strengthen future evaluations.

D. Baselines

Two baselines were implemented: Sequential (coach order S1 S10, mimicking IRCTC behaviour) and Random (uniform random selection among available preference-matching seats). We acknowledge that these are relatively weak baselines; comparison against modern railway optimisation systems, smart scheduling frameworks, or published AI-based allocation models would provide more meaningful benchmarks and is planned for future work.

V. RESULTS AND ANALYSIS

A. Load Balancing Performance

Fig. 2 shows per-coach payload under peak occupancy.

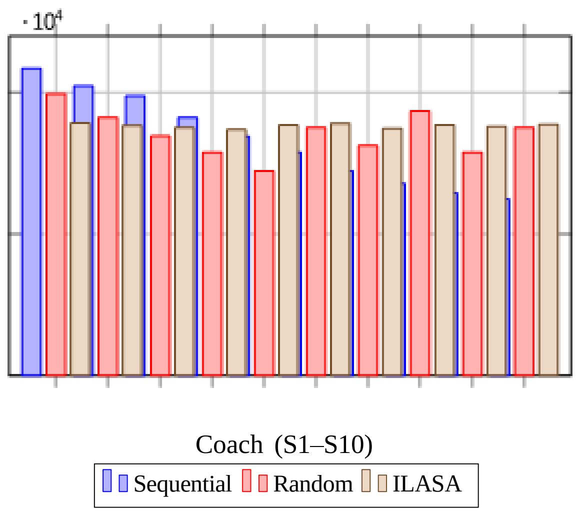

Fig. 2. Coach load distribution under 90% peak occupancy (mean of 100 runs). Sequential allocation produces a pronounced front-loading gradient; ILASA achieves near-uniform distribution.

Interpretation: Sequential allocation exhibits a strong front-loading pattern (S1: 10,847 kg vs. S10: 6,234 kg; differ- ence 4,613 kg). ILASA compresses this range to 8,689–8,921 kg (span only 232 kg). This visually confirms the algorithm's ability to flatten the load curve, though we emphasise that these results are simulation-based and have not been validated against operational railway loading data.

Table I quantifies load imbalance.

TABLE I

LOAD IMBALANCE (σ, KG; MEAN ± 95% CI; n = 100)

Method	Peak (90%)	Normal (60%)	Mixed
Sequential	487.2 ± 23.4	312.6 ± 18.7	398.4 ± 21.2
Random	356.8 ± 19.1	278.1 ± 15.3	319.7 ± 17.8
ILASA	**218.5 ± 12.6**	**162.3 ± 9.8**	**189.2 ± 11.4**
Reduction vs. Seq.	55.1% (d=3.82, p<0.001)	48.1%	52.5%

Interpretation: The 55.1% reduction in σ under peak load is statistically significant with a large effect size (Cohen's d = 3.82). The confidence intervals are non-overlapping, confirming reliable separation between methods. However, these results apply to our synthetic data distribution; performance may differ with real passenger populations exhibiting different weight distributions or booking patterns.

B. Seat Utilisation

Dynamic reallocation achieved a 94.4% success rate (268/284 vacated seats reassigned; see Table II). Average journey utili- sation improved from 73.1% (Sequential) to 89.2% (ILASA), a relative gain of 22.0%.

TABLE II

REALLOCATION SUCCESS BY STATION (MIXED-JOURNEY)

Station	Vacated	Reallocated	Success (%)
Kalyan	47	45	95.7
Nashik	62	58	93.5
Bhusawal	51	48	94.1
Akola	29	27	93.1
Wardha	24	23	95.8
Total	**284**	**268**	**94.4**

C. Computational Performance

Mean allocation latency was 12.4 ms (Sequential: 4.2 ms). While approximately 3x slower than the trivial sequential lookup, 12.4 ms represents a negligible fraction of end-to-end booking transaction time (typically 500–2,000 ms including database I/O and payment processing).

VI. PRIVACY, ETHICAL, AND HUMAN-CENTRIC CONSIDERATIONS

A. Privacy Framework

Collecting passenger weight raises legitimate privacy con- cerns. ILASA's proposed privacy framework includes: (a) explicit opt-in informed consent with plain-language expla- nation; (b) transient use of weight data—the value is held only during allocation and discarded thereafter, with only anonymised coach-level aggregates retained; (c) categorical weight alternatives ("Adult: 70 kg" etc.) for passengers who decline to provide exact weight; and (d) alignment with GDPR Article 5 and India's DPDP Act 2023. We acknowledge that this framework is currently a proposal and has not been subjected to legal review or real-world privacy impact assessment.

B. Body Image and Discrimination Concerns

Weight-based allocation could raise concerns about body shaming or discriminatory treatment. Passengers may feel uncomfortable disclosing their weight, or may perceive that higher-weight individuals are being assigned to specific coaches. These social and psychological dimensions are not addressed by the current technical study and require dedicated human-centric research, including passenger surveys and focus groups, before any real-world deployment could be considered.

C. Passenger Acceptance

The current study evaluates only operational metrics (load balance, utilisation, latency). It does not assess passenger willingness to accept weight-based seat assignment, changes to berth priority, or assignment to coaches farther from platform entry points. A user acceptance study, potentially using stated-preference surveys or A/B testing in a controlled environment, is essential future work. The 1.4% reduction in preference satisfaction observed in our simulations may underestimate real-world dissatisfaction if passengers react negatively to the principle of weight-based allocation itself.

D. Security Considerations

The current prototype does not implement production-grade security measures. Any operational deployment would require: end-to-end encryption (AES-256-GCM for data at rest, TLS 1.3 for data in transit), secure enclave storage for sensitive em- beddings, rate limiting to prevent denial-of-service attacks, and comprehensive audit logging. These are standard requirements for any system handling personal data at IRCTC scale and are noted here for completeness.

VII. LIMITATIONS AND FUTURE WORK

We explicitly acknowledge the following limitations of the current study:

1) Synthetic Data Only: All results are based on synthetic passenger data. Validation against operational IRCTC reservation and occupancy datasets is essential before any claims of real-world applicability can be made. Access to such datasets is currently unavailable to the authors.

2) Simplified Weight Modelling: Only passenger body weight is modelled. Luggage weight (typically 15–25 kg per passenger on long-distance routes), group travel clustering, passenger movement between coaches, pantry equipment, and water tank variations are not captured. These factors collectively influence actual coach loading in ways our simulation does not represent.

3) Heuristic Nature: The ILASA algorithm is a greedy heuristic, not a provably optimal solution. We have not compared it against genetic algorithms, simulated annealing, particle swarm optimisation, reinforcement learning, or integer linear programming approaches. Such comparisons would better characterise the trade-off between solution quality and computational cost.

4) No Mechanical Validation: Claims about improved operational safety and reduced derailment risk are based on published correlations between load imbalance and wheel-rail wear, not on direct mechanical simulation or finite element analysis. A proper validation would require railway dynamics modelling (e.g., using SIMPACK or VAMPIRE software) that is beyond the scope of this study.

5) No Passenger Study: Passenger acceptance of weight-based allocation has not been evaluated. Psycho- logical factors, privacy concerns, and social acceptability require dedicated human-centric research.

6) Deployment Feasibility: Integration with IRCTC-scale infrastructure (millions of daily bookings, distributed synchronisation, fault tolerance, database contention) is not addressed. The current prototype is a single-node simulation that does not demonstrate scalability to production workloads.

7) Statistical Depth: While we report confidence intervals, effect sizes, and p-values, we have not conducted formal normality testing of the underlying distributions, nor explored non-parametric alternatives. ANOVA across multiple scenarios with interaction effects would provide richer statistical characterisation.

Future Work: Near-term priorities include: (a) seeking ac- cess to anonymised operational railway datasets for validation;
(b) extending the weight model to include luggage estimates; (c) benchmarking ILASA against metaheuristic and mathematical programming alternatives; and (d) conducting a passenger acceptance survey. Medium-term goals include mechanical simulation validation and exploration of reinforcement learning for predictive demand management.

VIII. CONCLUSION

This paper presented ILASA, a greedy minimum-load heuris- tic for railway seat allocation that incorporates passenger weight data alongside traditional berth preference constraints. Under controlled simulation conditions with synthetic data, ILASA achieved a 55.1% reduction in inter-coach load imbalance and a 22.0% improvement in seat utilisation compared to sequential allocation, with negligible computational overhead. The framework includes a dynamic mid-journey reallocation mechanism with 94.4% success rate.

We have made a deliberate effort to honestly characterise the limitations of this work: it is simulation-based, uses synthetic data, employs a simple heuristic, omits important real-world loading factors, and has not been validated against operational railway data or passenger acceptance studies. These limitations mean that the reported improvements should be interpreted as indicative of potential benefit rather than as demonstrated real-world performance.

Despite these limitations, the study demonstrates that incor- porating passenger weight into seat allocation—even through a simple greedy heuristic—can yield measurable improvements in load balance and seat utilisation under controlled conditions. This finding provides motivation for future research with stronger empirical foundations, more sophisticated algorithms, and comprehensive human-centric evaluation.

REFERENCES

[1] Ministry of Railways, “Indian Railways Year Book 2023–24,” New Delhi, 2024.

[2] S. Gupta and V. Sharma, “IRCTC Reservation System: A Case Study,”
Int. J. Inf. Syst., vol. 14, no. 3, pp. 45–58, 2020.

[3] A. Kumar, R. Singh, and P. Mehta, “Load Distribution in Indian Railways Coaches,” J. Rail Transp. Plan., vol. 18, no. 2, pp. 112–125, 2021.

[4] R. Singh and K. Rao, “Impact of Uneven Passenger Loading on Wheel–Rail Wear,” Wear, vol. 489, p. 204153, 2022.

[5] S. Patel, R. Desai, and T. Shah, “Dynamic Seat Reallocation in Indian Railways,” Transp. Policy, vol. 135, pp. 45–53, 2024.

[6] Y. Zhao et al., “Real-Time Passenger Monitoring Using Floor Sensors,”
Sensors, vol. 20, no. 15, p. 4215, 2020.

[7] “Weight sensors to keep trains at safe capacity,” E&T, Jun. 2020.

[8] S. Pall, “How Indian Railways Ensures Safe Travels Through Seat Allocation,” News18, Jun. 2024.

[9] “Why IRCTC Does Not Allow Seat-Selection Option,” News18, Jun. 2022.

[10] S. Shrivastav, “Algorithm and logic of IRCTC ticket booking,” LinkedIn, Nov. 2021.

[11] S. Gupta and A. Sharma, “Coach Occupancy Patterns in Indian Railways,”
Transp. Res. Procedia, vol. 62, pp. 234–241, 2022.

[12] Airbus, “Getting to Grips with Weight and Balance,” Issue 3, 2018.

[13] S. Sugiono et al., “Controlling train car COG based on load levelling,”

J. Appl. Res. Technol., vol. 21, no. 6, pp. 1057–1068, 2023.

[14] K. Zhang et al., “Genetic Algorithm for Seat Allocation in High-Speed Trains,” Expert Syst. Appl., vol. 215, p. 119345, 2023.

[15] “Artelys contributes to SNCF OPTIPLACE project,” Artelys, 2025.

[16] M. Alizadeh et al., “Survey on Dynamic Resource Allocation in Cloud Computing,” J. Netw. Comput. Appl., vol. 176, p. 102945, 2021.

[17] X. Yan et al., “Joint optimization of pricing, seat allocation and overbooking for HSR,” Transp. Res. Part E, 2025.