EJSIT
European Journal of Science, Innovation and Technology
ISSN: 2786-4936
www.ejsit-journal.com
Volume 6 | Number 5 | 2026

A Framework for Diagnosing Underutilization in Distributed Vehicle Networks

Azamat Yerlanovich Sadykov*
Head of Industrial Automation
Engineering Development Center Kazakhstan
Almaty, Kazakhstan

ABSTRACT

Fleet operators diagnosing underutilization typically decompose it into demand-side insufficiency, supply-side excess, schedule misalignment, and maintenance clustering, four structural categories whose attribution rests on correlating utilization troughs against local demand density and vehicle-to-demand ratios. This article shows that the attribution procedure, built for fixed-route and scheduled logistics operations, reduces to two live branches once misalignment and maintenance clustering are settled separately, which leaves room to extend the framework to a documented condition in dispatch-mediated networks: aggregate demand and supply balanced within a zone-time cell while individual completion still falls short because vehicles and requests remain unpaired. A synthetic diagnostic exercise built on the same correlation logic the existing test uses confirms this directly: across 2,000 simulated replications, demand-side insufficiency and supply-side excess are each attributed correctly in 100% of cases, while a matching-friction condition of comparable trough depth produces correlation coefficients indistinguishable from zero on both channels and is attributed correctly in 0% of cases. Ride-hailing and taxi studies document the same pattern empirically, reporting completion shortfalls near forty percent even where local supply and demand intensities stay unremarkable. The article argues that distributed, dispatch-mediated vehicle networks benefit from a fifth attribution branch, matching friction, orthogonal to the four structural categories and applied at its own analytical level, so utilization diagnostics can move from route-based fleets to platform-coordinated ones while attributing the full observed deficit correctly.

Keywords: fleet utilization, underutilization diagnosis, matching friction, dispatch-mediated networks, vehicle rebalancing, utilization rate index

INTRODUCTION

A freight fleet operating ninety-two percent of its scheduled hours sits in the top quartile of conventional utilization benchmarks, yet once thirty-one percent of that operating time is absorbed by non-revenue repositioning and only eighty-five percent of dispatched pickups are actually completed, the fleet's effective productive yield falls to fifty-four percent of scheduled availability (Karazhekov, 2026). The gap between the headline figure and this productive reality reflects a real measurement problem: a single utilization ratio conflates the structurally distinct conditions that produce it, so an operator reading 0.92 must decompose the number first to establish whether the deficit sits in demand, in deployment, in scheduling, or in maintenance.

Distributed, dispatch-mediated networks, ride-hailing platforms and free-floating carsharing fleets chief among them, extend that picture, because a request and a vehicle can occupy the same zone-time cell in numbers the two zone-level ratios would call balanced, while the platform's matching layer still leaves a substantial share of them unpaired. Sections below

* Email: azamat.sadykov.ia@gmail.com

first establish, through the literature, that a pairing-stage effect with exactly this signature has been documented independently of fleet-sizing and positioning research, and then test directly, on simulated data built from the attribution framework's own decision rule, whether that rule captures the effect once it is present alongside a normal-looking demand density and vehicle-to-demand ratio.

LITERATURE REVIEW

Fielding, Glauthier, and Lave (1978) gave the field its first widely adopted productivity ratio, revenue vehicle-hours per employee, which quantified labor efficiency but counted every vehicle-hour of movement as equivalent regardless of whether that hour returned fare revenue. Karlaftis (2004) moved past a single ratio by applying Data Envelopment Analysis to vehicle-kilometers, passenger trips, and operating expenditure jointly, letting an efficiency frontier absorb several inputs and outputs at once; because that frontier is estimated at the system level, though, it returns a single system-level efficiency score that stays silent on any individual vehicle. De Borger, Kerstens, and Costa (2002) surveyed the resulting body of frontier studies and found that most left scale inefficiency entangled with allocative and technical inefficiency at the operational level, which is precisely the separation a fleet manager needs before deciding whether to resize the fleet, redeploy vehicles, or revise the schedule. Rodríguez González, Wilby, Vinagre Díaz, Fernández Pozo, and Sánchez Ávila (2023) tested this conflation with 449 days of minute-resolution car2go data across ten European cities, an empirical test of a claim the field had mostly argued theoretically, and showed that near-identical fleet-level rates can sit atop sharply different underlying vehicle-state distributions once the aggregate is disaggregated. Vinagre Díaz, Fernández Pozo, Rodríguez González, and Wilby (2024) pushed the same critique into e-scooter fleets, where measured utilization rates near 0.20 to 0.33 and out-of-service proportions above twenty-five percent coexisted with profitability close to one million euros in the Rome deployment studied, decoupling the diagnostic signal from the financial outcome it is usually assumed to track. Karazhekov (2026) responds to exactly this fragmentation, decomposing utilization into five independent components, revenue activity, non-revenue movement, idle efficiency, maintenance availability, and demand capture.

Fleet size and vehicle position turn out to be separable failure modes. George and Xia (2011) modeled a vehicle rental system as a closed queuing network and derived fleet size directly from a target service-availability level, establishing that steady-state utilization is a deterministic function of the ratio between fleet size and demand intensity once the network reaches equilibrium, a result that holds regardless of where within the network any given vehicle sits at a given moment. List, Wood, Nozick, Turnquist, Jones, Kjeldgaard, and Lawton (2003) extended this into demand uncertainty, formulating fleet sizing as a two-stage stochastic program that weighs the fixed cost of an additional vehicle against the recourse cost of insufficient capacity, treating every unit of oversupply as a quantifiable, avoidable cost, where earlier practice had absorbed it as an unavoidable buffer. Rahimi-Vahed, Crainic, Gendreau, and Rei (2015) showed that adding multi-depot and periodic routing constraints to a fleet-sizing heuristic changes the recommended fleet size relative to a single-depot formulation, and Pascual, Martinez, and Giesen (2013) found the same sensitivity when maintenance capacity is added as a joint constraint alongside fleet size in a cyclical transportation system; together these results matter for diagnosis because they mean a supply-side excess reading, measured only against demand, can misclassify a fleet that is correctly sized for demand but oversized relative to its own maintenance throughput. Xu, Meng, and Liu (2018) closed the loop between sizing and deployment by jointly optimizing electric vehicle fleet size alongside relocation and personnel assignment for one-way carsharing, showing that the two decisions interact: a fleet sized correctly in aggregate can still register local supply-side excess if the relocation policy leaves vehicles stranded away from where the aggregate calculation assumed they would serve

demand. Because sizing and positioning are therefore separable failure modes, each carrying its own deficit, a fleet sized correctly by any of the models above still leaves open whether its vehicles occupy the right zones at the right times.

Schmöller, Weikl, Müller, and Bogenberger (2015) documented exactly this in Munich and Berlin free-floating carsharing, showing idle time clustering spatially around low-demand zones even while the fleet's aggregate size matched citywide demand, which establishes that aggregating across zones hides a positioning failure in exactly the way aggregating across dimensions hides Fielding's conflation. Weikl and Bogenberger (2016) responded by building relocation plans from historical zone-level shortage and surplus indicators and Kypriadis, Pantziou, Konstantopoulos, and Gavalas (2020) formalized the relocation cost itself as an optimization objective; together these results show that measuring misallocation at the zone level makes it correctable through relocation. Yang, Wu, Sun, Qu, and Li (2021) extended relocation from historical averages to real-time requests, narrowing the condition under which the static plans of Weikl and Bogenberger hold: a static relocation plan stays adequate only while demand volatility within the relocation interval remains below the threshold at which real-time replanning would improve on it. Urban, Baur, Soppert, and Steinhardt (2025) complicate this group's implicit optimism about relocation, finding, from two years of operational data with a major European free-floating provider, that the supply-demand matching inefficiency underlying zone-level misallocation shrinks as system scale grows while a residual persists at every scale tested.

Xu, Sun, Liu, and Wang (2018) analyzed driver-response data from a Chinese ride-hailing platform and found that close to forty percent of hailing requests go unanswered by any driver, with response probability depending jointly on spatio-temporal supply and demand intensity together with economic incentives, beyond either intensity alone, an early sign that aggregate balance and individual pairing answer to different variables. Lu, Wang, Chen, and Xiong (2023) carried this further with large-scale sensor data, showing that a taxi driver's time efficiency, measured as the propensity to choose the fastest available route to a request, tracks the driver's accumulated experience and access to real-time routing intelligence independently of the surrounding demand and supply levels, so that two drivers facing an identical zone-time demand-supply ratio can still produce different completion outcomes purely from pairing-stage information asymmetry. The freight sector supplies a structural parallel: Bailey, Unnikrishnan, and Lin (2011) showed that minimizing backhaul costs through carrier collaboration depends on matching carriers holding empty capacity to shippers holding return loads, and McKinnon and Ge (2006) found empirically that only about two percent of UK food-sector empty truck-kilometers could be filled through backloading once trip-length and scheduling constraints were applied, a result that held even on lanes where aggregate freight volumes in both directions were roughly balanced. Bakir, Erera, and Savelsbergh (2022) built a robust rolling-horizon procedure for empty repositioning that explicitly weighs the cost of moving a vehicle empty against the cost of leaving it unavailable when demand later materializes, and ÇalIskan and Hall (2003) solved the equivalent problem for long-haul trucking networks with a dynamic integer programming model that returned feasible driver-home-time schedules on networks of up to thirty nodes and forty-eight periods using real carrier data; both treat the repositioning decision as a distinct optimization object from fleet sizing, confirming that a matching failure can coexist with balance in either aggregate quantity, standing apart from it as an independent condition.

METHODS

Whether the two-branch attribution rule registers a pairing-stage failure of the kind documented above is a question the reviewed studies leave open, since each ran the demand-density and vehicle-to-demand-ratio test on cases where the cause was already known to be structural. The exercise follows the same synthetic-validation logic Karazhekov (2026) uses to establish the FFT-based structural/stochastic distinction before applying it to operator data, and for the same reason: a diagnostic test should demonstrably recover a known cause before its output is trusted on data where the cause is unknown.

Three conditions were simulated independently over a 180-day horizon, a length chosen to exceed the ninety-day minimum the monograph reports as necessary to resolve monthly periodicities, so the test operates on a horizon long enough that its results stand independently of observation-window length. Each day, a zone-level demand density series was generated with a baseline mean of 100 requests, a weekly seasonal component of amplitude 10, and independent daily noise with standard deviation 11; a zone-level vehicle-to-demand ratio series was generated with baseline mean 1.0 and daily noise standard deviation 0.10; and a Utilization Rate Index (URI) series was generated with baseline mean 0.85 and daily noise standard deviation 0.06. A recurring ten-day trough window, occurring once every thirty days and covering one third of the horizon, was then imposed through exactly one of three channels depending on the condition: in the demand-side-insufficiency condition, demand density within the trough window was multiplied by 0.6 and the URI series was reduced by a term proportional to that shortfall (coefficient 0.22); in the supply-side-excess condition, the vehicle-to-demand ratio within the trough window was multiplied by 1.4 and URI was reduced by a term proportional to the ratio's elevation above 1.0 (coefficient 0.24); in the matching-friction condition, demand density and the ratio were left at their baseline distributions throughout, while URI was reduced within the trough window by a term proportional to the absolute value of an independent daily covariate representing dispatch-level information quality (coefficient 0.13), drawn from a standard normal distribution and, by construction, uncorrelated with either zone-level series. The three coefficients were set so that the resulting trough depth in URI, roughly 0.15 to 0.20 below baseline at the midpoint of the trough window, is comparable across conditions; a test that distinguishes the conditions correctly is therefore doing so on the basis of which covariate the trough tracks, holding trough severity itself constant.

For each of 2,000 independent replications per condition, the Pearson correlation was computed between the full 180-day URI series and the demand-density series, and separately between the URI series and the vehicle-to-demand-ratio series. Following the attribution logic set out in the Introduction, a replication was scored as attributed to demand-side insufficiency if the magnitude of the demand-density correlation reached 0.35 or higher, as attributed to supply-side excess if the magnitude of the ratio correlation reached 0.35 or higher, and as unattributed when both thresholds stayed unreached. The 0.35 cutoff was set at the lower bound of the moderate-correlation range, 0.35 to 0.55, that Karazhekov (2026) reports for the Revenue Activity/Idle Efficiency dimension pair in real transit and logistics data; adopting the lower, more permissive bound of that range as the attribution threshold gives the existing two-branch rule its best available chance of registering a signal. A stricter threshold would have biased the comparison toward the argument this article makes.

RESULTS

Table 1 reports the mean and standard deviation of both correlation coefficients, together with the resulting attribution rate, for each of the three simulated conditions. In the demand-side-insufficiency condition, the correlation between URI and demand density averaged 0.635 (SD = 0.041) across the 2,000 replications, while the correlation with the vehicle-to-demand ratio averaged 0.002 (SD = 0.075); every replication crossed the 0.35 threshold on the demand-density channel while the ratio channel stayed below it in every case, so the rule attributed 100% of cases to demand-side insufficiency and 0% to supply-side excess. The supply-side-excess condition produced the mirror result: a mean ratio correlation of -0.662 (SD = 0.040) against a mean demand-density correlation of 0.014 (SD = 0.070), with 100% of cases attributed to supply-side excess and 0% to demand-side insufficiency. Fig. 1 shows why the rule succeeds in both cases: in the top and middle panels, the colored covariate driving the trough (demand density in the first, the ratio in the second) visibly tracks the black URI line across the full horizon, which is exactly the co-movement a correlation-based test is designed to catch.

The matching-friction condition produced a starkly different pattern across the 2,000 replications. The mean correlation with demand density was 0.012 (SD = 0.071) and the mean correlation with the ratio was -0.002 (SD = 0.073), in both cases statistically indistinguishable from the correlation each structural condition produced on its own untouched channel (0.014 and 0.002, respectively, above), and every replication stayed below the 0.35 threshold on both channels; the attribution rate for this condition was 0% on both branches and 100% unattributed. The bottom panel of Fig. 1 shows the trough itself is fully present: the black URI line dips as deep and as often in the matching-friction panel as it does in the two panels above it, while the two colored covariates stay within their baseline noise band throughout, including during the trough windows.

Table 1. Mean correlation coefficients and attribution outcomes across 2,000 simulated replications per condition (attribution threshold |r| >= 0.35)

Condition	Mean r, URI-demand density (SD)	Mean r, URI-vehicle-to-demand ratio (SD)	Attributed: demand-side insufficiency	Attributed: supply-side excess	Unattributed
Demand-side insufficiency	0.635 (0.041)	0.002 (0.075)	100.0%	0.0%	0.0%
Supply-side excess	0.014 (0.070)	-0.662 (0.040)	0.0%	100.0%	0.0%
Matching friction	0.012 (0.071)	-0.002 (0.073)	0.0%	0.0%	100.0%

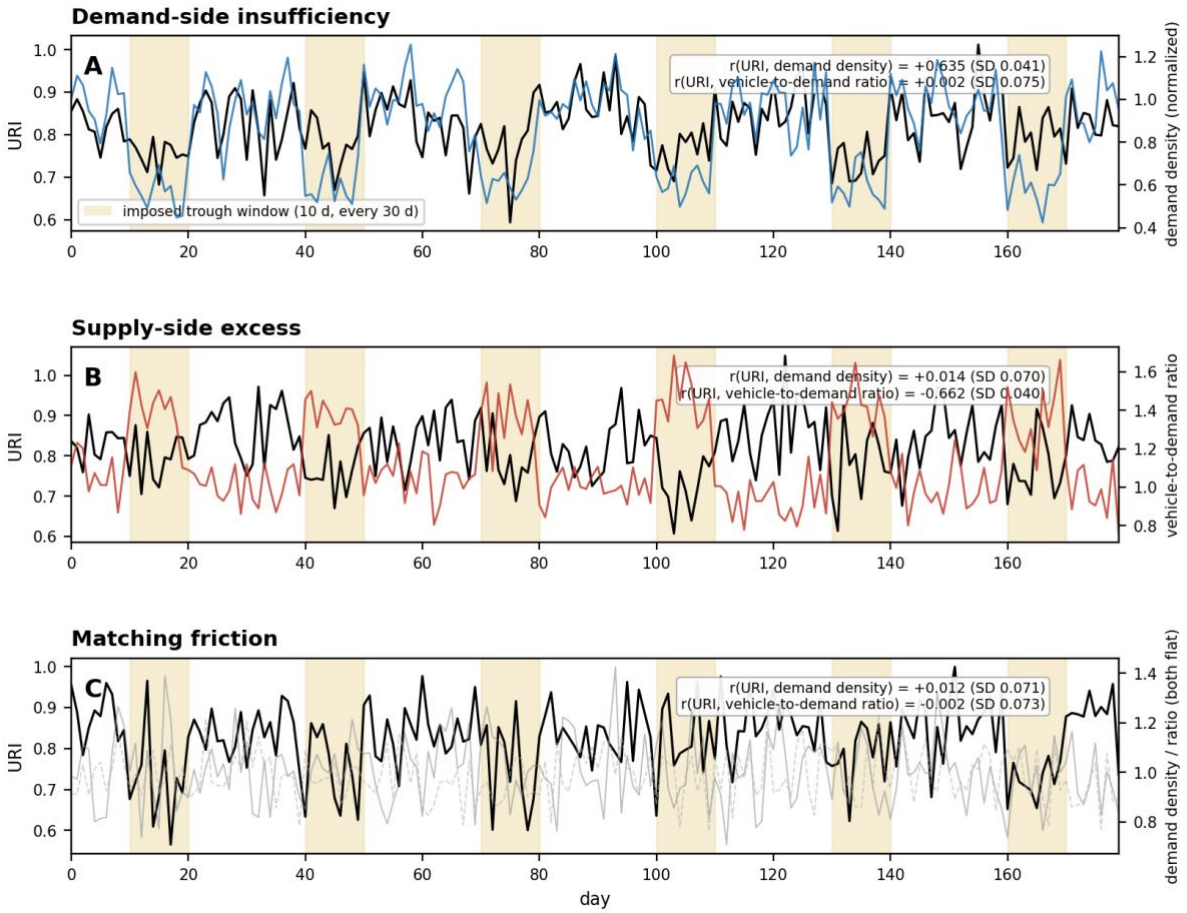

Fig. 1. Simulated URI trajectories under the three underutilization conditions (black: URI; colored: the zone-level covariate the attribution test checks)

DISCUSSION

The capacity-sizing and positioning literatures reviewed above give Karazhekov's (2026) structural taxonomy genuine external support: George and Xia's queuing result underwrites the supply-side excess category on independent theoretical grounds, and Weikl and Bogenberger's zone-level relocation triggers are functionally the same cross-correlation test the monograph proposes for attributing a utilization trough to demand-side insufficiency versus supply-side excess. Beyond the taxonomy itself, Karazhekov's (2026) closed-loop architecture carries the attribution output through a demand-forecasting stage and a prescriptive recommendation engine that quantifies projected URI impact and financial return before any intervention is authorized, a design credited with eight-to-fourteen-percent deadhead reductions and four-to-seven-percent revenue-vehicle-hour gains in transit deployments and ten-to-eighteen-percent idle-time reductions in logistics deployments, giving any additional attribution branch, including the matching-friction branch proposed here, a concrete path from diagnosis to operator action.

Karazhekov's (2026) attribution framework tests whether a persistent utilization trough traces to demand-side insufficiency, supply-side excess, schedule misalignment, or maintenance clustering, evaluating the first two against zone-level demand density and the vehicle-to-demand ratio. Both the simulation and the ride-hailing evidence point to a further condition worth adding to this framework, one in which both ratios stay unremarkable while completion still falls short. Response probability in Xu, Sun, Liu, and Wang (2018) depends jointly on spatio-temporal supply and demand intensity together with economic incentives, so the incentive term carries weight the local supply-demand ratio alone leaves unaccounted for,

which is the real-world counterpart of the independent incentive covariate that drove the simulated matching-friction condition to a 0.012 mean correlation with demand density, well below the 0.635 the demand-insufficiency condition produced on the same channel. Because the mechanism behind the shortfall in Lu, Wang, Chen, and Xiong's (2023) driver-level data sits in information, a driver's routing intelligence and accumulated experience determining whether a given request gets served, fleet size and zone boundary stay irrelevant to this outcome, and the freight parallel in McKinnon and Ge (2006) shows the same disconnect surviving in a setting where aggregate directional flows are balanced yet backloading still fails for scheduling and trip-length reasons that trace to timing and geometry, separate from volume entirely.

A cross-correlation procedure keyed only to demand density and the vehicle-to-demand ratio will, on encountering a matching-friction trough, either categorize it as the nearest of the four existing categories, most plausibly demand-side insufficiency, since a failed match can resemble a missed request from the demand side, or leave it as unattributed residual variance, exactly the outcome the simulation returned in 100% of matching-friction replications; in either case the intervention logic points toward redeployment or fleet resizing, whereas the underlying issue in this scenario sits in the dispatch layer itself. The extension this suggests takes the form of an orthogonal test, one that checks pairing outcomes conditional on demand and supply both sitting in normal range, standing apart from the existing four categories at a different analytical level entirely. Building on the monograph's cross-correlation step, a URI trough that stays weakly correlated with both demand density and the vehicle-to-demand ratio, the pattern the simulation shows occurring in 100% of matching-friction cases even at a lenient 0.35 threshold, calls first for a matching-friction test against completion rate, run conditional on driver-level or dispatch-level covariates, routing intelligence, accumulated experience, and incentive structure, the same covariates Xu et al. and Lu et al. identify as affecting completion independently of the local supply-demand ratio.

Because the matching-friction covariate was constructed to be exactly uncorrelated with demand density and the vehicle-to-demand ratio, the 0% attribution rate reported above shows that the two-branch rule reaches its boundary in the limiting case of a pairing failure carrying zero volumetric signature; the intermediate degrees of correlation stay outside this exercise's reach, where a partially demand-linked matching friction, for instance a platform whose incentive payouts scale with local demand intensity, might produce a weak, nonzero correlation with demand density and cross the 0.35 threshold by coincidence, independent of the true cause.

CONCLUSION

Diagnosing underutilization in a distributed vehicle network requires distinguishing five conditions in place of the four the field currently tests for: three that the accumulated fleet-sizing and rebalancing literature already instruments through demand-density and vehicle-to-demand-ratio testing, and that the simulation above recovers at a 100% attribution rate wherever the underlying literature identifies them (demand-side insufficiency, supply-side excess, schedule misalignment), a fourth, engineering-driven category (maintenance clustering) that sits outside the scope of the sources reviewed here, and a fifth, matching friction, that the ride-hailing and freight-matching literatures document and that the same simulated test attributes correctly in 0% of cases, because its defect sits in the pairing step, separate from either aggregate quantity the test checks. Extending Karazhekov's (2026) attribution procedure to dispatch-mediated networks means adding a conditional test, pairing outcome given normal-range demand and supply, positioned outside the existing decision tree at a level of its own, separate from a fifth branch added to the first two; the two tests answer different questions and belong in sequence.

For an operator building this kind of diagnostic system, the practical consequence is that a utilization trough uncorrelated with demand density or the vehicle-to-demand ratio, the exact signature the simulation attaches uniquely to matching friction among the conditions tested, counts as evidence to run the matching-friction test next, using the driver- and dispatch-level covariates that Xu, Sun, Liu, and Wang (2018) and Lu, Wang, Chen, and Xiong (2023) identify as affecting completion independently of the local supply-demand ratio. Because Urban, Baur, Soppert, and Steinhardt (2025) find this fifth category's contribution narrowing as a system scales, the practical urgency of building the matching-friction test is inversely related to system size, mattering most for mid-scale, dispatch-mediated fleets, regional ride-hailing markets and growing carsharing operators among them, a band too large for manual dispatch to absorb pairing failure and too small for scale to absorb it statistically.

REFERENCES

Bailey, E., Unnikrishnan, A., & Lin, D.-Y. (2011). Models for minimizing backhaul costs through freight collaboration. *Transportation Research Record: Journal of the Transportation Research Board, 2224*(1), 51-60. https://doi.org/10.3141/2224-07

Bakir, I., Erera, A., & Savelsbergh, M. (2022). A robust rolling horizon framework for empty repositioning. *Transportation Research Part C: Emerging Technologies, 144*, Article 103903. https://doi.org/10.1016/j.trc.2022.103903

ÇalIskan, C., & Hall, R. W. (2003). A dynamic empty equipment and crew allocation model for long-haul networks. *Transportation Research Part A: Policy and Practice, 37*(5), 405-418. https://doi.org/10.1016/S0965-8564(02)00033-2

De Borger, B., Kerstens, K., & Costa, Á. (2002). Public transit performance: What does one learn from frontier studies? *Transport Reviews, 22*(1), 1-38. https://doi.org/10.1080/01441640010020313

Fielding, G. J., Glauthier, R. E., & Lave, C. A. (1978). Performance indicators for transit management. *Transportation, 7*, 365-379. https://doi.org/10.1007/BF00168037

George, D. K., & Xia, C. H. (2011). Fleet-sizing and service availability for a vehicle rental system via closed queuing networks. *European Journal of Operational Research, 211*(1), 198-207. https://doi.org/10.1016/j.ejor.2010.12.015

Karazhekov, D. (2026). *Efficiency in transportation networks: Managing vehicles at scale*. GlobeEdit.

Karlaftis, M. G. (2004). A DEA approach for evaluating the efficiency and effectiveness of urban transit systems. *European Journal of Operational Research, 152*(2), 354-364. https://doi.org/10.1016/S0377-2217(03)00029-8

Kypriadis, D., Pantziou, G., Konstantopoulos, C., & Gavalas, D. (2020). Optimizing relocation cost in free-floating car-sharing systems. *IEEE Transactions on Intelligent Transportation Systems, 21*(9), 4017-4030. https://doi.org/10.1109/TITS.2020.2995197

List, G. F., Wood, B., Nozick, L. K., Turnquist, M. A., Jones, D. A., Kjeldgaard, E. A., & Lawton, C. R. (2003). Robust optimization for fleet planning under uncertainty. *Transportation Research Part E: Logistics and Transportation Review, 39*(3), 209-227. https://doi.org/10.1016/S1366-5545(02)00026-1

Lu, Y., Wang, Y., Chen, Y., & Xiong, Y. (2023). The role of data-based intelligence and experience on time efficiency of taxi drivers: An empirical investigation using large-scale sensor data. *Production and Operations Management, 32*(11), 3665-3682. https://doi.org/10.1111/poms.14056

McKinnon, A. C., & Ge, Y. (2006). The potential for reducing empty running by trucks: A retrospective analysis. *International Journal of Physical Distribution & Logistics Management, 36*(5), 391-410. https://doi.org/10.1108/09600030610676268

Pascual, R., Martinez, A., & Giesen, R. (2013). Joint optimization of fleet size and maintenance capacity in a fork-join cyclical transportation system. *Journal of the Operational Research Society, 64*(7), 982-994. https://doi.org/10.1057/jors.2012.103

Rahimi-Vahed, A., Crainic, T. G., Gendreau, M., & Rei, W. (2015). Fleet-sizing for multi-depot and periodic vehicle routing problems using a modular heuristic algorithm. *Computers & Operations Research, 53*, 9-23. https://doi.org/10.1016/j.cor.2014.07.004

Rodríguez González, A. B., Wilby, M. R., Vinagre Díaz, J. J., Fernández Pozo, R., & Sánchez Ávila, C. (2023). Utilization rate of the fleet: A novel performance metric for a novel shared mobility. *Transportation, 50*, 285-301. https://doi.org/10.1007/s11116-021-10244-x

Schmöller, S., Weikl, S., Müller, J., & Bogenberger, K. (2015). Empirical analysis of free-floating carsharing usage: The Munich and Berlin case. *Transportation Research Part C: Emerging Technologies, 56,* 34-51. https://doi.org/10.1016/j.trc.2015.03.008

Urban, L., Baur, A., Soppert, M., & Steinhardt, C. (2025). Scale matters: Why the efficiency of free-floating vehicle sharing systems depends on the system's scale. *Transportation Research Part C: Emerging Technologies, 178*. https://doi.org/10.1016/j.trc.2025.105169

Vinagre Díaz, J. J., Fernández Pozo, R., Rodríguez González, A. B., & Wilby, M. R. (2024). Time-based utilization rate of the fleet: Measuring deep inefficiencies in e-scooter services in Atlanta and Rome. *IEEE Transactions on Intelligent Transportation Systems, 25*(6), 4987-4997. https://doi.org/10.1109/TITS.2023.3330572

Weikl, S., & Bogenberger, K. (2016). Integrated relocation model for free-floating carsharing systems: field trial results. *Transportation Research Record: Journal of the Transportation Research Board, 2536*(1), 19-27. https://doi.org/10.3141/2536-03

Xu, K., Sun, L., Liu, J., & Wang, H. (2018). An empirical investigation of taxi driver response behavior to ride-hailing requests: A spatio-temporal perspective. *PLOS ONE, 13*(6), e0198605. https://doi.org/10.1371/journal.pone.0198605

Xu, M., Meng, Q., & Liu, Z. (2018). Electric vehicle fleet size and trip pricing for one-way carsharing services considering vehicle relocation and personnel assignment. *Transportation Research Part B: Methodological, 111*, 60-82. https://doi.org/10.1016/j.trb.2018.03.001

Yang, S., Wu, J., Sun, H., Qu, Y., & Li, T. (2021). Double-balanced relocation optimization of one-way car-sharing system with real-time requests. *Transportation Research Part C: Emerging Technologies, 125*. https://doi.org/10.1016/j.trc.2021.103071