Austrian Journal of Technical
and Natural Sciences, 2026 No 7–8
ISSN 2310-5607

PREMIER
Publishing
ppublishing.org

Section 6. Transport

DOI:10.29013/AJT-26-7.8-129-134

THE HIDDEN COST LAYER: DEADHEAD, UTILIZATION THRESHOLDS, AND MARGIN FORMATION IN FREIGHT CARRIER OPERATIONS

Shevchenko Maxim Igorevich [1]

[1] Head of the Transport Analytics Laboratory, Institute of Applied Transportation Management Systems, Kazan

Cite: *Shevchenko M.I. (2026). The Hidden Cost Layer: Deadhead, Utilization Thresholds, and Margin Formation in Freight Carrier Operations. Austrian Journal of Technical and Natural Sciences 2026, No 7–8.* https://doi.org/10.29013/AJT-26-7.8-129-134

Abstract

Carrier operating cost reported as a single per-mile figure conceals a compositional structure that determines profitability. This paper synthesizes econometric, technical, and carrier-economics literatures around one organizing question: how does a reduction in deadhead mileage or an improvement in fleet utilization translate into margin, and does that translation depend on where the affected assets sit relative to a break-even threshold? The analysis applies a cost-decomposition and threshold-translation method built on published literature. It separates operating cost into a visible layer and a submerged layer comprising detention, deadhead, and administrative friction, and it evaluates routing, network, and procurement interventions according to whether they plausibly move a fleet segment across its break-even utilization floor. The synthesis shows that algorithmic routing advances set a ceiling on deadhead reduction; carrier scale, network topology, and procurement structure set a floor beneath which no algorithmic gain can compensate for structural disadvantage.

Keywords: *deadhead reduction; fleet utilization; break-even threshold; submerged cost; carrier margin; vehicle routing; freight carrier economics*

Introduction

Operating cost reported as a single per-mile figure conceals a compositional structure that determines whether a freight carrier is profitable. Treating that figure as a flat ledger obscures exactly the variation that determines margin. A narrow set of cost drivers accounts for most of the variance in carrier margin. The remainder splits between two layers: a visible one that appears directly on profit-and-loss statements, and a submerged one that erodes margin without registering as a discrete line item. Doskach (2026 b) describes this distinction as a cost iceberg: detention time, deadhead repositioning, and administrative friction together account for a further 14 to 22 percent of effective operating cost once the opportunity cost of idle capital equipment is included. Doskach develops this cost-iceberg framework from direct observation of carrier operations, drawing on the kind of granular detail that rarely surfaces in industry-facing reporting.

Doskach's (2026 a) break-even utilization floor is the minimum weekly loaded-mile threshold below which a given truck generates negative net revenue. Below that floor, submerged cost stops being a marginal drag and starts determining whether the asset is profitable at all. Absorbing a deadhead mile costs a carrier operating comfortably above the floor only reduced profit. The same mile costs a carrier operating near or below it the difference between solvency and loss on that unit of equipment. Crossing the floor changes an asset's profitability status. That threshold structure is the object this paper analyzes.

Econometric, technical, and carrier-economics literatures each address a piece of this problem without connecting them. Econometric studies of empty miles quantify the scale of the deadhead problem and its structural causes, yet stop short of translating a percentage-point reduction into a margin effect that depends on where an affected vehicle sat relative to a profitability threshold. Vehicle-routing research has produced increasingly capable algorithms for reducing deadhead and improving dispatch efficiency. Those algorithms are evaluated almost exclusively against distance or cost objectives, disconnected from the threshold-crossing logic that determines whether an efficiency gain changes a carrier's solvency. Carrier-economics research on scale, procurement, and network structure explains why margin varies across carriers and segments, without incorporating the submerged-cost decomposition that identifies where that variance originates operationally.

Materials and Methods

This paper applies a cost-decomposition and threshold-translation method built on the published literature synthesized below. A submerged-cost taxonomy, introduced above, supplies the decomposition mechanism, separating operating cost into visible line items and a submerged layer comprising detention, deadhead, and administrative friction that behaves differently from the visible layer under threshold-crossing analysis. Doskach's (2026 a) break-even utilization floor supplies the translation mechanism: a claimed efficiency improvement, whether from routing technology, network redesign, or procurement strategy, is evaluated according to whether it plausibly moves the fleet segment it addresses across a profitability threshold, not according to its raw magnitude alone.

This paper draws evidence from four bodies of literature, moving from the scale of the deadhead problem to the mechanisms available for reducing it, the carrier-level structures that determine how much of a reduction actually reaches margin, and the exogenous shocks that test whether that architecture holds under stress. Econometric studies of empty-mile incidence and its structural causes establish the scale of the deadhead problem and the mechanisms generating it (Yang, 2024; Heilmann, 2020; Allen et al., 2024). Technical operations-research literature on vehicle routing establishes the algorithmic solution space available for deadhead reduction and its performance characteristics under realistic demand uncertainty (Nazari et al., 2018; Kool et al., 2019; Ulmer et al., 2020). Carrier-economics literature on scale, network structure, and procurement establishes how margin varies across carrier configurations independent of any single routing intervention (Miller & Muir, 2020; Muir et al., 2019; Acocella & Caplice, 2023; Acocella et al., 2020; Owczarek, 2024; Rodrigue, 2020; Solakivi & Ojala, 2024; Izadi et al., 2019). Regulatory and technology-shock literature establishes how an exogenous change to the operating environment stresses margin architecture and reveals which carriers were operating close to a profitability threshold before the shock (Scott et al., 2021; Miller et al., 2021; Trick et al., 2021; Miller, Muir, et al., 2020).

This design carries limitations that follow directly from its structure, and the most consequential concerns the evidentiary status of the cost-decomposition figures themselves. The percentages attached to the submerged-cost taxonomy applied here originate in a practitioner synthesis (Doskach, 2026 b) and have not been independently replicated against the carrier-level panel data used elsewhere in this paper. The threshold logic built on this taxonomy treats the figures as illustrative of an order of magnitude documented in industry practice; the argument holds regardless of whether any single percentage is exactly correct.

Results

Fuel, driver wages, equipment depreciation, and insurance dominate visible operating cost, and carriers that manage exclusively against this visible total systematically misallocate improvement effort. Izadi et al. (2019), reviewing cost models used across the road freight literature, found that the majority of published models categorize cost by broad functional type, such as operational, value-of-time, and external cost, without a structural mechanism for identifying which categories respond nonlinearly to utilization change; this gap is precisely what Doskach's (2026 b) submerged-cost taxonomy addresses, separating fixed, variable, and quasi-fixed cost components and showing that the quasi-fixed category generates discontinuous margin cliffs as utilization crosses replacement-cycle thresholds, a step pattern that conventional per-mile accounting cannot represent.

A carrier whose utilization sits just below a quasi-fixed threshold, such as a tire replacement cycle or a maintenance interval, absorbs the full fixed component of that cost over a smaller loaded-mile base than a carrier whose utilization crosses the threshold. The distinction changes which carriers stay solvent. Two carriers with identical average deadhead percentages can face materially different margin outcomes depending on which side of a quasi-fixed threshold their fleets predominantly occupy, a distinction that a single aggregate deadhead statistic cannot reveal and that motivates the threshold-based framework this paper applies throughout.

The practical consequence extends to how carriers allocate improvement budgets. A capital investment in equipment-focused efficiency, such as an aerodynamic retrofit, produces a bounded and well-documented improvement on a cost category that is already visible and already accounted for in most fleet budgets. A comparable investment directed at dispatch-layer coordination targets the submerged category directly, and because that category is measured poorly or not at all in conventional reporting, its baseline is frequently understated and its achievable improvement correspondingly underestimated by managers comparing the two investment options using visible-cost metrics alone. This measurement asymmetry drives the persistent industry pattern of overinvestment in equipment specification relative to dispatch coordination. The underlying economic opportunity in the two categories is the same.

The mechanisms generating empty miles have been isolated with increasing precision at the level of individual carrier decisions. Heilmann (2020), studying how a digital freight brokerage platform introduced an algorithmic load-bundling feature, estimated a 22.6 percent reduction in deadhead miles following adoption, equivalent to an increase in measured truck utilization from 81.6 to 85.8 percent, and attributed the improvement specifically to resolving information frictions that had previously prevented drivers from identifying compatible return loads within their available operating window. Allen et al. (2024), documenting trip-chaining behavior among Colombian truckers, found that drivers frequently construct complex multi-stop sequences instead of executing simple out-and-back trips, avoiding empty repositioning by design, evidence that deadhead avoidance behavior is already a first-order consideration in driver decision-making, and that further gains depend on improving the information available to support that behavior.

Vehicle routing research has approached deadhead reduction primarily as an algorithmic optimization problem. Two methodological shifts over the past decade show how far that toolkit has moved. Nazari et al. (2018) demonstrated that a reinforcement-learning policy trained through policy-gradient methods could match or exceed classical routing heuristics on stochastic capacitated instances without hand-engineered rules specific to a given network structure. Kool et al. (2019) extended this approach with an attention-based encoder-decoder architecture, showing that self-attention mechanisms generalize across routing problem variants more effectively than recurrent alternatives, a property directly relevant to carriers whose lane structure changes as their customer base evolves.

Ulmer et al. (2020) formalized the distinction between static and dynamic routing architectures, which determines which of these algorithmic advances actually reduce deadhead in practice. Policies that incorporate anticipated future demand information systematically outperform reactive dispatch rules under conditions of meaningful demand uncertainty, and the performance gap between the two widens as uncertainty increases. Static planning performs within a narrow margin of dynamic optimization under low volatility, but that gap grows large enough under high volatility to determine whether a fleet segment operates above or below a utilization threshold. This threshold sensitivity connects Ulmer's routing finding directly to the threshold framework applied throughout this paper: an algorithmic improvement that appears modest when measured in average deadhead percentage can be consequential when it disproportionately benefits vehicles operating near the break-even floor.

The value of anticipatory routing varies sharply across a carrier's network, a point the aggregate performance statistics in this literature tend to obscure. Because static planning already approximates the dynamic optimum under low volatility, lanes with stable, predictable demand gain comparatively little from a shift toward forecast-integrated dispatch; the disproportionate gains sit instead on lanes with volatile, thin, or seasonally concentrated demand, precisely where reactive dispatch diverges furthest from an anticipatory policy. A carrier deploying dynamic routing technology uniformly across its network, without segmenting lanes by volatility, captures only a fraction of the available benefit on its most valuable lanes and still pays the full implementation cost across lanes where the technology adds little.

Network topology sets a further ceiling on how much of a routing algorithm's theoretical gain a carrier can actually realize, because it determines whether deadhead is reabsorbed or simply relocated. Point-to-point structures generate deadhead at the destination of every unmatched load, since no return network exists to reabsorb it; Rodrigue (2020) contrasts this with hub-and-spoke topologies, which absorb the same deadhead through network-level backhaul balancing instead of letting it accumulate at the endpoint. Equipment choice interacts with this topology constraint instead of substituting for it: Solakivi and Ojala (2024), comparing high-capacity trucks against standard semitrailers, found that larger vehicles are not uniformly more cost-efficient, since their advantage depends on achieving a higher load factor than the equipment they replace, which confirms that utilization determines realized margin more than equipment specification alone does.

Procurement structure determines how much of a utilization gain survives contact with the market a carrier actually sells into. Acocella and Caplice (2023), reviewing the truckload procurement literature comprehensively, identified a persistent gap between strategic procurement decisions and execution-stage outcomes, with committed capacity frequently going unused or unavailable capacity being sourced at spot-market premiums that erode the margin a routing improvement would otherwise deliver. This execution-stage gap varies across carriers: Acocella et al. (2020) found that carrier behavior in response to shipper tender patterns varies systematically between carriers that behave reciprocally, accepting a stable volume in exchange for stable future tenders, and those that behave opportunistically, which is why the margin benefit of a utilization improvement depends on the shipper relationship structure within which it operates as much as on the improvement itself. Because relationship structure mediates how much of a utilization gain a carrier retains, the pricing mechanism through which that gain converts into revenue matters just as much: Owczarek (2024), modeling less-than-truckload pricing for light commercial vehicles, showed that the intensity of vehicle usage and the number of loading and unloading points jointly determine achievable pricing, confirming at the carrier level that utilization intensity functions as a pricing variable in its own right, beyond its role as a cost input.

These procurement-side findings interact with the threshold framework in a way that is easy to overlook. A carrier whose fleet sits near its break-even floor is also the carrier with the least slack to absorb the spot-market premiums that Acocella and Caplice (2023) identify as a common consequence of procurement mismatch. The same undercapitalization that pushes a fleet toward its utilization threshold in the first place also limits that fleet's ability to negotiate the kind of reciprocal shipper relationship that Acocella et al. (2020) find reduces opportunistic exposure, so the carriers most exposed to threshold-crossing risk from deadhead are frequently the same carriers least equipped to manage procurement-side margin risk through relationship structure alone.

Discussion

The evidence assembled here supports a layered view of utilization, whose components determine margin through different mechanisms instead of through one performance percentage. Doskach's (2026b) distinction between the visible and submerged cost layers explains why carriers that manage exclusively against reported per-mile cost systematically underinvest in dispatch-layer improvement relative to equipment-focused interventions with smaller theoretical ceilings but greater accounting visibility. This visible/submerged distinction functions as a practical prioritization tool: it identifies which specific lever, dispatch-layer coordination or equipment-focused investment, a given carrier should prioritize given its position relative to the break-even floor, instead of leaving that choice to whichever category happens to be easier to measure.

The routing literature and the carrier-economics literature synthesized in this paper answer different questions that are frequently conflated in practice. Nazari et al. (2018), Kool et al. (2019), and Ulmer et al. (2020) establish that algorithmic dispatch improvement is technically achievable and that its value scales with demand uncertainty, a finding about the ceiling on deadhead reduction; Miller and Muir (2020), Muir et al. (2019), and Rodrigue (2020) show that the same routing improvement produces different margin outcomes depending on carrier scale, specialization, and network topology, which sets the floor beneath which no algorithmic improvement can compensate for structural disadvantage. A carrier evaluating a routing technology investment needs both answers, and treating the algorithmic ceiling as sufficient evidence of margin recovery, without reference to the carrier's position relative to its break-even floor, overstates the return that investment will deliver.

The regulatory and turnover evidence reviewed in this paper indicates that margin architecture is tested most clearly during shocks, when carriers already close to their break-even thresholds are pushed across them by a uniform external change. Scott et al. (2021) and Miller et al. (2021) show that a uniform compliance mandate produces uneven consequences precisely because carriers entered the shock at different positions relative to their break-even thresholds, a pattern consistent with the threshold-crossing logic this paper applies to deadhead reduction.

Conclusion

The threshold-translation method applied throughout this paper changes what an efficiency claim has to answer. A percentage-point reduction in deadhead, or a percentage-point gain in utilization, carries no fixed margin value on its own: its value depends on where the affected fleet segment sat relative to its break-even floor before the intervention. Routing technology sets a ceiling on how much deadhead a carrier can eliminate; carrier scale, network topology, and procurement structure set the floor beneath which no amount of algorithmic improvement closes the remaining gap. Regulatory mandates and driver turnover reveal this floor instead of creating it: they push already-exposed carriers across a line that was there before the shock.

References

Acocella, A., & Caplice, C. (2023). Research on truckload transportation procurement: A review, framework, and future research agenda. Journal of Business Logistics, – 44(2). – P. 228–256. URL: https://doi.org/10.1111/jbl.12333

Acocella, A., Caplice, C., & Sheffi, Y. (2020). Elephants or goldfish? An empirical analysis of carrier reciprocity in dynamic freight markets. Transportation Research Part E: Logistics and Transportation Review, 142, Article 102073. URL: https://doi.org/10.1016/j.tre.2020.102073

Allen, T., Atkin, D., Cantillo Cleves, S., & Hernández, C. E. (2024). The traveling trucker problem. AEA Papers and Proceedings, – 114. – P. 334–339. URL: https://doi.org/10.1257/pandp.20241044

Doskach, H. (2026 a). Building a profitable trucking company in the U.S. LAP LAMBERT Academic Publishing.

Doskach, H. (2026 b). Logistics efficiency and cost reduction strategies. LAP LAMBERT Academic Publishing.

Heilmann, K. (2020). Information frictions, load matching, and route efficiency in the trucking industry (SSRN Working Paper No. 3545019). URL: https://doi.org/10.2139/ssrn.3545019

Izadi, A., Nabipour, M., & Titidezh, O. (2019). Cost models and cost factors of road freight transportation: A literature review and model structure. Fuzzy Information and Engineering, – 11(3). – P. 257–278. URL: https://doi.org/10.1080/16168658.2019.1706960

Kool, W., van Hoof, H., & Welling, M. (2019). Attention, learn to solve routing problems! International Conference on Learning Representations (ICLR 2019).

Miller, J. W., & Muir, W. A. (2020). A new perspective on returns to scale for truckload motor carriers. Journal of Business Logistics, – 41(3). – P. 236–258.

Miller, J. W., Muir, W. A., Bolumole, Y., & Griffis, S. E. (2020). The effect of truckload driver turnover on truckload freight pricing. Journal of Business Logistics, – 41(4). – P. 294–309. URL: https://doi.org/10.1111/jbl.12252

Miller, J. W., Scott, A., & Williams, B. (2021). Pricing dynamics in the truckload sector: The moderating role of the electronic logging device mandate. Journal of Business Logistics, – 42(4). – P. 388–405. URL: https://doi.org/10.1111/jbl.12256

Muir, W. A., Miller, J. W., Griffis, S. E., Bolumole, Y. A., & Schwieterman, M. A. (2019). Strategic purity and efficiency in the motor carrier industry: A multiyear panel investigation. Journal of Business Logistics, – 40(3). – P. 204–228.

Nazari, M., Oroojlooyjadid, A., Snyder, L. V., & Takáč, M. (2018). Reinforcement learning for solving the vehicle routing problem. Advances in Neural Information Processing Systems, – 31. – P. 9839–9849.

Owczarek, P. (2024). Price valuation modeling of less-than-truckload (LTL) shipments for financial continuity assurance. European Research Studies Journal, – 27(3). – P. 209–224.

Rodrigue, J.-P. (2020). The geography of transport systems (5th ed.). Routledge.

Scott, A., Balthrop, A., & Miller, J. W. (2021). Unintended responses to IT-enabled monitoring: The case of the electronic logging device mandate. Journal of Operations Management, – 67(2). – P. 152–181. URL: https://doi.org/10.1002/joom.1110

Solakivi, T., & Ojala, L. (2024). Environmental and economic potential of high-capacity trucks. Case Studies on Transport Policy. URL: https://doi.org/10.1016/j.cstp.2024.101284

Trick, S., Peoples, J., & Ross, A. (2021). Driver turnover in the trucking industry: What's the cost of reducing driver quit rates? Research in Transportation Economics, 89, Article 101129. URL: https://doi.org/10.1016/j.retrec.2021.101129

Ulmer, M. W., Goodson, J. C., Mattfeld, D. C., & Thomas, B. W. (2020). On modeling stochastic dynamic vehicle routing problems. EURO Journal on Transportation and Logistics, 9, Article 100008. URL: https://doi.org/10.1016/j.ejtl.2020.100008

Yang, R. (2024). Geographic imbalance, search frictions, and regulation: Causes of empty miles in freight trucking (World Bank Policy Research Working Paper No. 10775).

submitted 15.06.2026;
accepted for publication 29.06.2026;
published 31.08.2026

Contact: shevchenkomaksig@gmail.com