



DOI:10.29013/AJT-26-5.6-283-290



DEMAND-COUPLED FLEET LIFECYCLE OPTIMIZATION: A JOINT FRAMEWORK FOR PRICING, MAINTENANCE SCHEDULING, AND DISPOSITION TIMING

Vitalii Kolesnykov¹

¹ K Motors In, Owner, USA

Cite: Kolesnykov, V. (2026). *Demand-Coupled Fleet Lifecycle Optimization: a Joint Framework for Pricing, Maintenance Scheduling, and Disposition Timing*. *Austrian Journal of Technical and Natural Sciences* 2026, No 5–6. <https://doi.org/10.29013/AJT-26-5.6-283-290>

Abstract

Fleet lifecycle management in commercial vehicle rental operations is characterized by a structural decoupling of pricing, maintenance scheduling, and disposition decisions that systematically transfers mileage-intensive bookings to mechanically degraded vehicles, accelerating depreciation beyond the threshold at which residual value recovery remains commercially viable. This paper presents a demand-coupled fleet lifecycle optimization framework that treats vehicle physical condition as an endogenous continuous state variable jointly constraining pricing feasibility bounds, maintenance scheduling cost, and disposition timing within a scalarized dual-objective function solved by alternating finite-horizon dynamic programming and closed-form concave pricing optimization. The framework is evaluated on a 420-vehicle urban rental fleet over a 12-month operational horizon. Key results include an 18.7% reduction in maintenance scheduling opportunity cost relative to calendar-based scheduling, a 6.1 percentage-point improvement in residual value realization rate relative to the fixed 36-month calendar replacement policy.

Keywords: *fleet lifecycle optimization, dynamic pricing, maintenance scheduling, vehicle disposition, condition-based optimization, residual value modeling, two-stage decomposition, demand forecasting, rental fleet management, total cost of ownership*

1. Introduction

Unplanned mechanical failure in commercial rental fleets costs between 2.4 and 3.8 times the equivalent scheduled service event (Jardine & Tsang, 2006), and the dominant operational response to this fact has been to improve maintenance scheduling in isolation, refining condition thresholds, tightening service windows, adopting prog-

nostic health monitoring, without addressing the upstream commercial decision that determines how much mechanical stress a vehicle accumulates per day. Booking assignment in rental operations is governed by revenue management systems that optimize price and availability from demand forecasts without access to vehicle condition state, producing a feedback structure in which vehicles

in the lower quartile of mechanical condition, already approaching the cost-escalation phase of their service life, receive no pricing or booking-type differentiation from newly inducted units of the same class. Hartman & Tan (2014) document that across a century of equipment replacement models, utilization intensity has been treated as an exogenous parameter rather than a co-optimized decision variable, an assumption defensible for capital equipment on fixed operating schedules but structurally incorrect for rental assets whose mileage accrual is determined by commercial booking decisions executed within the same operational cycle as maintenance and disposition.

This paper introduces a demand-coupled fleet lifecycle optimization framework that resolves this structural gap by treating vehicle physical condition as an endogenous continuous state variable jointly constraining all three operational decision domains within a single scalarized objective solved by a two-stage decomposition alternating finite-horizon dynamic programming and closed-form concave pricing optimization.

2. Literature Review

The classical economic life framework optimizes replacement timing against a given utilization trajectory, which means any sensitivity result it produces, such as Figliozzi & Boudart's (2014) finding that moderate mileage input changes shift optimal bus replacement timing by two to three years, is a sensitivity to an assumption rather than a co-optimized output. Hartman & Murphy (2006) and Keles & Hartman (2004) minimize equivalent annual cost as a function of calendar age and cumulative operating cost, and Hartman & Tan (2014) extend the framework to variable utilization and technological change, but in every variant the mileage accrual rate enters the model as a scenario parameter set before optimization begins.

Stochastic dynamic programming applied to fleet resource allocation (Godfrey & Powell, 2002a, 2002b) defines vehicle state by location and availability and assigns vehicles to tasks over a rolling horizon using adaptive nonlinear value function approximations that produce near-optimal solutions on deterministic problems and substantial-

ly outperform rolling-horizon heuristics on stochastic ones. The shared assumption is operational interchangeability within a class, which is valid for logistics fleet balancing, where the question is which vehicle goes to which location, but eliminates from the model the condition dimension that determines whether assigning a specific vehicle to a specific booking profile is commercially rational. Turan et al. (2020) partially relax interchangeability by tracking individual asset condition trajectories in a simulation-based renewal model under uncertainty, but renewal timing in their formulation informs the replacement trigger without feeding back into in-period pricing or booking assignment; the condition trajectory and the commercial decision stream remain decoupled.

Condition-based maintenance scheduling models treat remaining useful life estimates as inputs to scheduling cost functions, and the limitation they share is the treatment of rental revenue as an exogenous scalar. Ostermeier & Henke (2021) incorporate RUL predictions explicitly into a multi-objective evolutionary algorithm over a rolling horizon, a methodological advance over calendar-interval scheduling, but the revenue opportunity cost term is a mean fleet utilization rate, not a period-specific demand forecast. A maintenance event scheduled during a regional demand peak incurs opportunity costs several multiples above the mean-utilization estimate, and the model has no mechanism to identify or avoid those dates.

3. Method


Each vehicle v at planning period t is assigned a lifecycle score $L(v, t) \in [0, 100]$ by the predictive lifecycle modeling engine, computed as a weighted linear combination of six sub-indices: mechanical condition, brake and tire wear, powertrain thermal stress, suspension load, mileage-adjusted age, and maintenance compliance. Sub-index weights are vehicle-class-specific, estimated monthly by penalized ordinary least squares regression against composite outcome scores derived from a 24-month disposition cohort in which the dependent variable combines observed maintenance costs and realized resale values in a historically calibrated ratio. The residual value trajectory projects wholesale

resale value at future disposition dates using a double-exponential base depreciation function of projected calendar age and projected odometer reading, corrected by a gradient-boosted ensemble term taking as inputs the current lifecycle score, the first derivative of cumulative maintenance cost over the preceding 60 days, and the deviation of the current external residual value benchmark from the base model's class-specific prior estimate.

Three categories of decision variables are defined over a planning horizon of T discrete daily periods: a base daily rental price for each vehicle class and period, a maintenance service date for each vehicle approaching a service threshold, and a recommended disposition period for each vehicle. The joint objective $\max \lambda \cdot R_{\text{fleet}} - (1 - \lambda) \cdot D_{\text{fleet}} + \omega \cdot V_{\text{terminal}}$ governs fleet-wide rental revenue R_{fleet} , physical depreciation cost D_{fleet} , and terminal residual value recovery V_{terminal} , with

$\lambda \in [0.10, 0.90]$ as the revenue-lifecycle trade-off coefficient and $\omega \in [0.05, 0.50]$ as the residual value weight. The depreciation term incorporates a monotonically increasing wear-rate amplification function of lifecycle score, capturing the acceleration of mechanical degradation per kilometer driven as vehicle condition declines (Jardine & Tsang, 2006). For vehicles with $L(v, t)$ below the utilization threshold of 42/100, the pricing ceiling is scaled by a continuous degradation discount function reaching 0.78 at the threshold value and decreasing toward lower scores, a formulation replacing the binary step-function threshold of the initial specification, whose discontinuity at 42/100 breaks the concavity of the second-stage pricing sub-problem, and the booking-type filter excludes reservations exceeding three days or 500 km.

Figure 1. System Architecture Block Diagram



The first stage of the two-stage decomposition solves vehicle-level maintenance scheduling and disposition timing independently per vehicle by finite-horizon dynamic programming with backward induction, using a state space defined by lifecycle score and cumulative maintenance cost, with terminal value set to the projected residual value at the recommended disposition date. The maintenance scheduling cost function has two terms: expected revenue loss from vehicle unavailability during service, computed as the product of the demand forecast for the candidate period, the workshop capacity utilization index, and the vehicle's trailing 30-day mean daily revenue; and the probability-weighted cost differential between an unplanned breakdown event and an equivalent scheduled service, using fault-category-specific ratios, 3.8 for drivetrain events, 3.1 for brake system, 1.2 for tire replacement, 1.1 for fluid service, rather than the fleet-wide flat constant of 2.4 that the initial specification carried. The second stage solves fleet-level pricing with maintenance and disposition decisions fixed from the first stage; separability across vehicle classes and periods reduces each pricing sub-problem to closed-form maximization of a concave revenue function evaluated at its first-order optimality condition. The stages alternate until the relative change in joint objective value between successive iterations falls below 0.01%.

The primary evaluation dataset covers 12 months of operational records from a 420-vehicle urban rental fleet across five geographic zones and five vehicle classes: B-segment compact, C/D-segment mid-size sedan, B-SUV compact sport utility, D-segment full-size sedan, and C-MPV compact minivan. The dataset contains 94,380 daily vehicle-state records from onboard telematics, 18,640 rental transaction records, 1,240 maintenance event records, and 52 weekly residual value benchmark updates from an external automotive valuation service. The 24 months preceding the evaluation period supply the disposition cohort for lifecycle weight estimation and the training dataset for demand forecasting and residual value modeling. Fleet size at evaluation start was 420 vehicles, varying between 388 and 431

over the horizon due to dispositions and acquisitions.

Convergence behavior and wall-clock computation time were additionally evaluated on two synthetic fleet instances, 50 vehicles across two zones and two vehicle classes, and 3,000 vehicles across ten zones and five vehicle classes, generated by resampling the empirical dataset with replacement and scaling demand parameters proportionally to fleet size, each run for 30 independent daily optimization cycles.

The calendar scheduling baseline assigns maintenance events according to manufacturer-recommended calendar intervals without demand sensitivity. The Revenue-Only pricing baseline maximizes rental revenue without lifecycle-state feedback, applying uniform class-level pricing to all vehicles regardless of condition. Both baselines were applied to the same 420-vehicle dataset over the same 12-month evaluation horizon under identical geographic zone structure, vehicle class composition, demand environment, and workshop capacity constraints. Maintenance opportunity cost reduction results are valid only for service windows of at least 8 calendar days; windows narrower than 8 days eliminate the demand-trough identification benefit. Residual value realization rate results hold for replacement vehicle lead times below 89 days.

4. Results

Across all fleet configurations in the primary evaluation dataset, the two-stage decomposition satisfied the 0.01% joint objective tolerance within four iterations, a convergence speed that does not degrade monotonically with fleet size. On the 420-vehicle fleet, the full daily optimization cycle completed in 4.7 minutes of wall-clock time; at the 3,000-vehicle synthetic fleet, wall-clock time reached 11.4 minutes, with the dynamic programming stage accounting for 74% of total compute time at both scales. The near-linear scaling behavior follows from the first-stage separability across individual vehicles: dynamic programming sub-problems are solved in parallel per vehicle, so the marginal cost of adding a vehicle to the fleet is the cost of one additional DP instance rather than a global state space expansion. The bi-

nary step-function threshold implementation fails this scaling property: tested as an alternative to the continuous degradation formulation across the 30 synthetic daily cycles at the 3,000-vehicle scale, it produced convergence failure in 6.3% of cycles, with iteration counts of 9–12 concentrated in fleet configurations containing a high proportion of vehicles with lifecycle scores between 39 and 45, the range straddling the 42/100 discontinuity, confirming that the concavity break introduced by the step function, not fleet size, is the binding convergence constraint.

Across the 1,240 maintenance events in the primary evaluation dataset, the demand-integrated scheduling model reduced mean revenue opportunity cost per maintenance event by 18.7% relative to the calendar scheduling baseline, with both models applied to the identical 420-vehicle fleet under the same workshop capacity ceiling of 15% per class-zone and identical manufacturer-prescribed service windows. The reduction mechanism is date selection within the permissible ser-

vice window guided by the demand forecast and workshop capacity utilization index: events are placed on dates where the demand forecast registers a local minimum and the capacity utilization index falls below 0.6, minimizing the product of these two factors with the vehicle's trailing mean daily revenue. Service windows narrower than 8 calendar days, present in 9.1% of the evaluated events, concentrated in vehicles within 200 km of the maximum mileage extension limit, eliminate the demand-trough identification benefit, reducing the opportunity cost advantage to 2.3% and making the model's output statistically indistinguishable from calendar scheduling over that sub-population.

Table 1 reports residual value realization rates for three vehicle classes under the proposed framework at the breakeven disposition trigger ($\theta = 0$) and under the fixed 36-month calendar replacement policy, evaluated on the 214 vehicles disposed of during the evaluation period against realized auction and dealer-channel resale prices.

Table 1. *Residual Value Realization Rate by Vehicle Class and Disposition Method*

Vehicle Class	Fixed 36-Month Rule	Proposed Frame- work ($\theta = 0$)	Difference (pp)	Mean Disposition Age, Proposed (months)
B-segment compact	0.841	0.893	+5.2	33.1
C/D-segment mid-size sedan	0.819	0.882	+6.3	31.7
B-SUV compact	0.803	0.871	+6.8	30.4
Fleet mean	0.821	0.882	+6.1	31.7

The 4.3-month mean early exit relative to the 36-month calendar rule captures residual value by exiting vehicles before the inflection point at which the depreciation curve steepens, a timing advantage that narrows to 4.4 percentage points after correcting for the 47-day mean replacement vehicle lead time, because the continuation indicator's net revenue projection does not deduct the vacancy cost during the procurement interval. At a lead time of 89 days, the corrected realization advantage reaches zero, and the fixed 36-month rule outperforms the framework's breakeven trigger. The B-SUV class, carrying the largest raw realization advantage at 6.8 percentage points and the earliest mean

disposition age at 30.4 months, experiences the sharpest lead-time sensitivity, with mean procurement lead times of 58 days in the evaluation fleet reducing the corrected advantage to 4.1 percentage points.

Increasing λ from the default 0.65 to 0.90 raises projected annual fleet-wide rental revenue by 8.3% and increases physical depreciation cost by 14.1%, a revenue-to-depreciation amplification factor of 1.70 that worsens to 2.31 at $\lambda = 0.88$. Setting $\lambda = 0.35$ reduces depreciation cost by 11.2% at a revenue cost of 9.6%. The residual value weight ω is a second-order sensitivity: varying ω across [0.05, 0.50] with λ fixed at 0.65 shifts mean fleet disposition age by 3.8 months

and changes annual resale recovery by 7.1%, but sensitivity of the scheduling and pricing outputs to ω remains below 0.3% across the full range, confirming that ω and λ govern orthogonal decision dimensions within the framework. The condition-tier elasticity limitation overstates revenue for sub-threshold vehicles by 4.7% at $\delta = 0.78$ and 8.2% at $\delta = 0.50$, inflating the continuation indicator's net revenue component and delaying the disposition recommendation for vehicles in the (20, 42) score range at precisely the operating phase where per-kilometer depreciation is highest.

5. Discussion

The 18.7% maintenance scheduling opportunity cost reduction, measured against the calendar scheduling baseline across 1,240 events on the 420-vehicle evaluation fleet, is almost certainly a conservative estimate of the gap between demand-integrated and demand-agnostic scheduling, and it is conservative for a reason that reflects poorly on the prior literature's reported improvements. The 18.7% figure is computed against a calendar baseline; demand-agnostic models that use mean utilization rates understate opportunity costs by construction, so their self-reported improvements are directionally inflated and not directly comparable.

The residual value findings confirm the mechanism identified by Kleindorfer et al. (2012) that early exit from fleet service captures value by preceding the steepest depreciation segment, and extend it from fleet-level renewal timing to individual vehicle condition-state-sensitive disposition. The 4.3-month mean early exit is consistent with the finite-horizon replacement analysis of Hartman & Murphy (2006), where optimal replacement timing consistently precedes the infinite-horizon economic life estimate, a convergence that holds here because the continuation indicator's 12-month projection horizon is structurally analogous to a finite planning horizon in which the terminal value is the recoverable residual, not an assumed replacement chain. The lead-time correction that reduces the realization advantage from 6.1 to 4.4 percentage points, and eliminates it at 89 days, is a structural gap in the continuation indicator formulation that no pub-

lished fleet renewal model, including the simulation-based approach of Turan et al. (2020) and the risk-averse renewal models of Ansariipoor et al. (2016), explicitly corrects for. Fleet operators in electrification transitions face procurement lead times that routinely exceed 60 days for battery electric vehicles, as documented by Juang et al. (2024), placing those deployments in the range where the lead-time gap materially reduces the framework's realization advantage below the 4.4 percentage-point corrected figure.

The amplification factor of 1.70 between revenue gain and depreciation cost increase above $\lambda = 0.65$, worsening to 2.31 at $\lambda = 0.88$, contradicts the approximately linear Pareto frontier that standard multi-objective fleet renewal models present near default operating points (Ansariipoor et al., 2016; Turan et al., 2020). The discrepancy arises from the convexity of the wear-rate amplification function in declining lifecycle score, a relationship that models without vehicle-level condition state as an endogenous variable cannot represent. As the optimizer pushes revenue at high λ values, it progressively assigns more intensive booking profiles to vehicles in the 42–55 score range, where the booking filter does not apply but where accelerated degradation accumulates; because the amplification function is convex, each successive unit of revenue extracted from this sub-population costs more in depreciation than the previous unit, producing the non-linear worsening of the trade-off ratio that flat-frontier models cannot predict.

The condition-tier elasticity bias, 4.7% revenue overstatement at the default degradation discount of $\delta = 0.78$, rising to 8.2% at $\delta = 0.50$, does not merely reduce the framework's precision at low lifecycle scores; it inverts the framework's operational logic for a specific and bounded sub-population. The framework's central claim is that coupling vehicle condition state to pricing, scheduling, and disposition produces better asset recovery outcomes than treating these domains independently. For vehicles with lifecycle scores in (20, 42), the coupling runs through the continuation indicator, which computes the disposition trigger as the discounted difference between projected net revenue and projected residual value loss. An 8.2%

overstatement of net revenue for the most degraded vehicles shifts the continuation indicator upward by the same margin, directly suppressing the disposition recommendation for the sub-population that is simultaneously accumulating the highest per-kilometer depreciation and generating the most inflated revenue projections. The coupling that is the framework's strength for vehicles above 42/100 becomes a systematic disposition-delay mechanism for vehicles below it, and the delay is largest at $\delta = 0.50$, the configuration an operator would choose when maximizing the intensity of condition-sensitive pricing, precisely the operating regime in which the framework is most aggressively applied. Correcting this requires condition-tier-stratified own-price elasticity estimation: four quartile-specific elasticity estimates rather than one class-level estimate, demanding within-tier sample sizes sufficient for double-logarithmic regression stability.

6. Conclusion

Coupling vehicle condition state to pricing, maintenance scheduling, and disposi-

tion within a single daily optimization cycle makes the revenue-depreciation trade-off nonlinearity operationally visible for the first time: the marginal depreciation cost of incremental revenue gain is not constant but amplifies to a factor of 2.31 per unit of revenue at $\lambda = 0.88$, a structural relationship that no prior model, lacking vehicle-level condition as an endogenous state variable, could detect, let alone optimize against. The question the framework cannot yet answer is whether the minimum-trace hierarchical demand forecast reconciliation retains calibration accuracy during rapid fleet composition changes in which the historical transaction distribution underpinning the temporal fusion transformer shifts structurally faster than the weekly retraining cycle corrects, a condition that applies directly to mixed-technology electrification transitions and would corrupt the period-specific opportunity cost rankings that determine which maintenance scheduling assignments differ from calendar scheduling and which do not.

References

- Ansaripoor, A. H., Oliveira, F. S., & Liret, A. (2016). Recursive expected conditional value at risk in the fleet renewal problem with alternative fuel vehicles. *Transportation Research Part C: Emerging Technologies*, – 65. – P. 156–171. URL: <https://doi.org/10.1016/j.trc.2015.12.010>
- Figliozzi, M. A., & Boudart, J. (2014). Vehicle technologies and bus fleet replacement optimization: Problem properties and sensitivity analysis utilizing real-world data. *Public Transport*, – 6(1). – P. 25–50. URL: <https://doi.org/10.1007/s12469-014-0086-z>
- Godfrey, G. A., & Powell, W. B. (2002a). An adaptive dynamic programming algorithm for dynamic fleet management, I: Single period travel times. *Transportation Science*, – 36(1). – P. 21–39. URL: <https://doi.org/10.1287/trsc.36.1.21.570>
- Godfrey, G. A., & Powell, W. B. (2002b). An adaptive dynamic programming algorithm for dynamic fleet management, II: Multiperiod travel times. *Transportation Science*, – 36(1). – P. 40–54. URL: <https://doi.org/10.1287/trsc.36.1.40.572>
- Hartman, J. C., & Murphy, A. (2006). Finite-horizon equipment replacement analysis. *IIE Transactions*, – 38(5). – P. 409–419. URL: <https://doi.org/10.1080/07408170500380054>
- Hartman, J. C., & Tan, C. H. (2014). Equipment replacement analysis: A literature review and directions for future research. *The Engineering Economist*, – 59(2). – P. 136–153. URL: <https://doi.org/10.1080/0013791X.2013.862891>
- Jardine, A. K. S., & Tsang, A. H. C. (2006). *Maintenance, replacement, and reliability: Theory and applications*. CRC Press / Taylor & Francis. ISBN: 9780849339660.
- Juang, J., Williams, W. G., Ramshankar, A. T., Schmidt, J., Xuan, K., & Bozeman, J. F. (2024). A multi-scale lifecycle and technoeconomic framework for higher education fleet electrification. *Scientific Reports*, 14, Article 4788. URL: <https://doi.org/10.1038/s41598-024-54752-z>

- Keles, P., & Hartman, J. C. (2004). Case study: Bus fleet replacement. *The Engineering Economist*, – 49(3). – P. 253–278. URL: <https://doi.org/10.1080/00137910490498951>
- Kleindorfer, P. R., Neboian, A., Roset, A., & Spinler, S. (2012). Fleet renewal with electric vehicles at La Poste. *Interfaces*, – 42(5). – P. 465–477. URL: <https://doi.org/10.1287/inte.1120.0640>
- Ostermeier, M., & Henke, M. (2021). Optimizing the maintenance schedule for a vehicle fleet: A simulation-based case study. *Engineering Optimization*, – 54(8). – P. 1–20. URL: <https://doi.org/10.1080/0305215X.2021.1919888>
- Turan, H. H., Elsayah, S., & Ryan, M. J. (2020). A long-term fleet renewal problem under uncertainty: A simulation-based optimization approach. *Expert Systems with Applications*, – 145, Article 113158. URL: <https://doi.org/10.1016/j.eswa.2019.113158>

submitted 20.05.2026;
accepted for publication 04.06.2026;
published 30.06.2026
© Kolesnykov V.
Contact: kolesnykovvitaliicase@gmail.com