

Threshold Coincidence and the Conditional Value of Lifecycle Decision Architecture in Rental Fleet Management

Jason Miller *

Senior Fleet Operations Manager NorthStar Fleet Services Inc., Denver, USA

*Correspondence: Jason Miller (jmiller.northstarfleet@gmail.com)

Abstract: Lifecycle architectures for rental fleets couple pricing, vehicle assignment, maintenance scheduling and disposal timing, and the profit they add over revenue-based scoring has two sources that behave differently. Mileage steering changes the odometer reading at which a vehicle reaches disposal, and decision-focused forecast training reduces ranking regret at that point. Both sources matter only where the cost-optimal disposal mileage lies within a band around an odometer mark at which wholesale prices drop. The band half-width equals the residual-value forecast error divided by the slope of the retention margin. On realized retirements of midsize sedans at one enterprise rental operator, 28.2% of vehicles fall inside the band. Coupling adds \$206 to \$228 per vehicle there and \$7 to \$13 elsewhere, differences indistinguishable from zero. Mileage steering supplies 60.7% of the premium and survives a step-aware forecaster; the decision-focused component falls from \$67.60 to \$8.20 once the forecaster learns the price drops. Richer telematics raise profit by \$86 to \$118 per vehicle in every cell, so architecture and data are separable contributions. Operators can restrict the cost of coupling to vehicles inside the band.

Keywords: Rental Fleet Management; Lifecycle Decision Architecture; Disposal Timing; Residual Value; Decision-Focused Learning; Predictive Maintenance

How to cite this paper:

Miller, J. (2026). Threshold Coincidence and the Conditional Value of Lifecycle Decision Architecture in Rental Fleet Management. *Universal Journal of Business and Management*, 5(1), 31-37. DOI: 10.31586/ujbm.2026.6837

Received: August 6, 2026

Revised: September 17, 2026

Accepted: September 28, 2026

Published: September 30, 2026



Copyright: © 2026 by the author. Submitted for possible open-access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Wholesale auction data on more than 22 million used-car transactions show price drops of about \$150 to \$200 at each 10,000-mile odometer mark from 10,000 to 100,000 miles, with smaller drops at 1,000-mile marks (Lacetera et al., 2012) [1]. A car sold at 79,900 to 79,999 miles fetched on average about \$210 more than one at 80,000 to 80,100 miles, and only about \$10 less than one at 79,800 to 79,899 miles (Lacetera et al., 2012) [1].

Sellers react to the marks: auction volume spikes just before each 10,000-mile threshold, and the evidence attributes the price pattern primarily to inattention among final buyers of used cars (Lacetera et al., 2012) [1]. Busse et al. (2013) [2] estimated the salience effect in wholesale and retail car markets alike. For a rental operator the pattern turns the last few hundred miles of a vehicle's life into a priced decision, because replacement theory already ties disposal timing to resale value. With several parallel assets, optimal replacement depends on age and cumulative utilization, salvage values respond to the utilization pattern, and a decision maker who allocates workload controls that pattern, with a threshold policy optimal under common cost assumptions (Hartman, 2004) [3]. A price function that jumps at round mileages makes workload allocation a lever on the sale price.

Rental operators hold that lever because the assignment of vehicles to reservations belongs to them. Oliveira et al. (2017) [4] name fleet and decision-making flexibility as the properties that distinguish car rental within revenue management. Steinhardt and Gönsch (2012) [5] integrate capacity control with planned upgrades, which assigns reservations

to vehicle classes, and Oliveira et al. (2019) [6] join pricing to fleet capacity under stochastic demand. Pricing and assignment are therefore the decisions through which mileage accrues, and a policy that uses them to aim mileage at a price mark needs a residual-value model whose shape matches the wholesale price function.

Revenue scoring for retention and disposal timing (Kolesnykov, 2026b) [7] times a vehicle's exit from its expected revenue. Any such score needs a residual-value term, since salvage value governs replacement timing (Hartman, 2004) [3], and a term that varies smoothly with mileage misprices the vehicle at each mark by an amount of the order of the jump, which is \$150 to \$200 on the auction scale above.

A policy is sequential when it fixes disposal timing from such a score and treats pricing and maintenance as downstream tasks, and coupled when it chooses the three jointly. A vehicle is threshold coincident when its cost-optimal disposal point lies close enough to a price discontinuity for a smooth forecast to reverse the retain-or-dispose ranking; the forecast error and the slope of the retention margin set that distance. Two questions follow. Does the advantage of coupling concentrate on threshold-coincident vehicles? And how much of an advantage reported for a lifecycle architecture belongs to the architecture, when the same deployments also bring richer sensor data? Answering both takes an evaluation in which decision architecture and data richness vary independently on vehicles retired near price marks.

2. Methods

The evaluation uses realized lifecycle records from an enterprise rental operator with regional airport and suburban stations. The dataset covers vehicles bought new and retired through physical and digital wholesale auctions between 2021 and 2025. To remove the influence of macroeconomic swings in residual values, the sample is restricted to the main passenger class, standard midsize sedans with four-cylinder gasoline engines, retired at odometer readings between 65,000 and 95,000 miles. Each vehicle record holds daily telematics summaries (cumulative mileage, engine hours, diagnostic trouble codes and sensor-based brake pad wear estimates), localized demand curves, realized rental rates, scheduled maintenance line items and the wholesale hammer price.

Balanced cohorts are built by nearest-neighbor propensity score matching, so that policy architecture is separated from fleet composition. The covariates are acquisition model year, months in service, acquisition cost and average utilization intensity before the final 10,000 miles of service. Vehicles are assigned to evaluation arms at 60,000 miles, the planning horizon that precedes the first terminal disposal boundary.

Wholesale price realization follows the discontinuous drops at round odometer marks documented by Lacetera et al. (2012) [1]. Let P_i be the transaction price of vehicle i and x_i its terminal odometer reading. The jump magnitudes come from a semi-parametric regression discontinuity specification:

$$\ln(P_i) = \alpha + \beta_1 x_i + \sum_{k \in K} \gamma_k \mathbb{1}\{x_i \geq x_k^*\} + \sum_{j \in M} \delta_j \mathbb{1}\{x_i \geq z_j^*\} + X_i \theta + \varepsilon_i \quad (1)$$

Here $K = \{70000, 80000, 90000\}$ lists the major 10,000-mile thresholds, and $M = \{66000, 67000, \dots, 94000\} \setminus K$ indexes the minor 1,000-mile marks. The vector X_i holds a cosmetic condition score from 1 to 5, season of sale, auction venue fixed effects and the day of week of the transaction. The coefficient γ_k identifies the log-price jump at threshold x_k^* , and J_k denotes the corresponding jump in dollars.

Replacement theory under stochastic utilization (Hartman, 2004) [3] ties disposal timing to the retention margin. For vehicle i at odometer reading x , the margin $m_i(x)$ is the expected operating profit of keeping the vehicle for one more 30-day decision period against immediate liquidation:

$$m_i(x) = E[R_i(\Delta x)] - E[C_i(x, \Delta x)] - E[V_i(x) - V_i(x + \Delta x)] \quad (2)$$

Here R_i is the rental revenue earned over the additional mileage Δx , C_i the direct operating and scheduled maintenance cost, and $V_i(x)$ the net salvage value function. The cost-optimal disposal point $x_{d,i}$ solves $m_i(x_{d,i}) = 0$.

Because m_i crosses zero at $x_{d,i}$, an error e in the salvage value moves the computed disposal point by e divided by the absolute slope of m_i at the crossing, to first order. The slope $|m'_i(x_{d,i})|$ is computed numerically by local linear regression of observed net monthly margins on cumulative mileage within 5,000 miles on either side of $x_{d,i}$. For a residual value model with mean absolute out-of-sample error e , the coincidence half-width is

$$w_i = \frac{e}{|m'_i(x_{d,i})|} \quad (3)$$

A vehicle is threshold coincident when its counterfactual optimal disposal mileage lies within the half-width of a major threshold, that is, when the following condition holds for at least one $k \in K$:

$$|x_{d,i} - x_k^*| \leq w_i \quad (4)$$

A 2×2 factorial evaluation crosses decision architecture with data richness (Table 1).

Table 1. Factorial arms crossing decision architecture with data richness

	Baseline telematics (odometer only)	Enriched telematics (sensor stream)
Sequential policy	Arm 1	Arm 2
Coupled policy	Arm 3	Arm 4

Arm 1, sequential with baseline telematics. Revenue scoring sets disposal timing from smoothed residual value estimates (Kolesnykov, 2026b) [7]. Pricing follows standard demand-curve clearance rules, and maintenance follows fixed OEM distance intervals.

Arm 2, sequential with enriched telematics. Disposal scoring adds high-dimensional telematics health indices from stacked ensembles (Kolesnykov, 2026c; Theissler et al., 2021) [8,9]. Pricing and vehicle assignment remain independent downstream operations.

Arm 3, coupled with baseline telematics. Dynamic pricing, assignment and disposal timing are optimized jointly (Kolesnykov, 2026a) [10], with calendar age and odometer reading as the only state information. A rolling-horizon dynamic program steers mileage toward the preferred liquidation windows.

Arm 4, coupled with enriched telematics. The full lifecycle architecture combines joint pricing, predictive maintenance scheduling and salvage timing, conditioned on calibrated telematics failure probabilities and the auction jump marks. Maintenance timing follows the usage pattern (de Jonge & Jakobsons, 2018) [11].

All arms share one pricing elasticity module calibrated on market booking curves, which keeps the gains of pricing optimization out of the architectural comparison (Besbes & Zeevi, 2009) [12]. A further sequential specification replaces the smooth residual forecaster with a step-aware model, gradient boosted trees with explicit jump indicators, and tests whether a forecaster that has learned the price drops reproduces the coupling premium.

Inside the coincidence band, the profit advantage of coupling over sequential revenue scoring splits into three components. Let Π_a denote mean profit per coincident vehicle in Arm a . Then

$$\Delta\Pi_{\text{coupled}} = \Pi_4 - \Pi_2 = \Delta\Pi_{\text{steer}} + \Delta\Pi_{\text{dl}} + \Delta\Pi_{\text{int}} \quad (5)$$

$$\Delta\Pi_{\text{int}} = (\Pi_4 - \Pi_3) - (\Pi_2 - \Pi_1) \quad (6)$$

The steering gain $\Delta\pi_{\text{steer}}$ is the additional salvage value obtained by altering reservation assignments so that the vehicle leaves at $x_k^* - \varepsilon$ in place of $x_k^* + \varepsilon$. The decision-loss gain $\Delta\pi_{\text{dl}}$ is the reduction in ranking regret from training residual forecasters against the downstream disposal objective (Elmachtoub & Grigas, 2022) [13]. The interaction term $\Delta\pi_{\text{int}}$ is the extra coupled advantage that enriched telematics bring at the boundary, the value of sensor-based condition monitoring (Prytz et al., 2015; Rögnvaldsson et al., 2018) [14,15] when the policy is able to act on it. Equations (5) and (6) imply that $\Delta\pi_{\text{steer}} + \Delta\pi_{\text{dl}} = \pi_3 - \pi_1$.

3. Results

The estimates show price penalties at each major mark: \$174.20 at 70,000 miles, \$218.60 at 80,000 miles and \$191.40 at 90,000 miles (Table 2). The average of about \$195 falls inside the range of \$150 to \$200 that Lacetera et al. (2012) [1] report for wholesale auctions. The local margin slope steepens with cumulative mileage, from \$0.182 per mile at 70,000 miles to \$0.264 at 90,000 miles, consistent with growing maintenance risk and falling rental rates. Base forecast errors of \$248.50, \$262.10 and \$281.00 then give half-widths of 1,365, 1,219 and 1,064 miles, so the band contracts as mileage rises. The share of retirements inside the band is 27.4% at 70,000 miles, 29.8% at 80,000 miles and 24.1% at 90,000 miles, and 28.2% of all retirements qualify as threshold coincident.

Table 2. Estimated price jumps and coincidence bands at the major odometer marks

Threshold x_k^*	Estimated jump J_k	Margin slope $ m'(x_d) $	Base forecast error e	Half-width w	Fleet share inside band
70,000 miles	\$174.20 $\pm$ \$18.40	\$0.182 per mile	\$248.50	1,365 miles	27.4%
80,000 miles	\$218.60 $\pm$ \$21.10	\$0.215 per mile	\$262.10	1,219 miles	29.8%
90,000 miles	\$191.40 $\pm$ \$19.80	\$0.264 per mile	\$281.00	1,064 miles	24.1%

Note. Values after $\pm$ are standard errors. Base forecast error is the mean absolute out-of-sample error of the smooth residual value model.

Table 3 gives net profit per vehicle over the final 15,000 operating miles for each arm, split by coincidence status. Outside the band, coupling changes profit by \$7 per vehicle under baseline telematics and by \$13 under enriched telematics, and neither difference departs from zero ($p > 0.10$). Inside the band, coupling adds \$206 per vehicle under baseline telematics ($t = 4.82$, $p < 0.001$) and \$228 under enriched telematics ($t = 5.16$, $p < 0.001$). Enriched telematics raise profit in every cell, by \$86 and \$92 per vehicle outside the band and by \$96 and \$118 inside it.

Table 3. Net profit per vehicle over the final 15,000 operating miles, by arm and coincidence status

Policy architecture	Telematics depth	Mean profit, coincident $ x_d - x_k^* \leq w$	Mean profit, non- coincident $ x_d - x_k^* > w$	Full cohort weighted mean
Arm 1: Sequential	Baseline (odometer)	\$3,412 $\pm$ \$42	\$3,895 $\pm$ \$35	\$3,759 $\pm$ \$28
Arm 2: Sequential	Enriched (sensors)	\$3,508 $\pm$ \$39	\$3,981 $\pm$ \$33	\$3,848 $\pm$ \$26
Arm 3: Coupled	Baseline (odometer)	\$3,618 $\pm$ \$41	\$3,902 $\pm$ \$36	\$3,822 $\pm$ \$27
Arm 4: Coupled	Enriched (sensors)	\$3,736 $\pm$ \$38	\$3,994 $\pm$ \$34	\$3,921 $\pm$ \$25

Mileage steering supplies \$138.40 of the \$228.00 premium on coincident vehicles, 60.7% of the total (Table 4). Coupled policies route coincident vehicles to shorter airport-transfer reservations and adjust rates to slow mileage accrual over the final 1,500 miles. In Arm 4, 84.1% of coincident vehicles were liquidated at odometer readings between $x_k^* - 400$ and $x_k^* - 50$ miles, which keeps them clear of the price drop. Under the sequential policies of

Arms 1 and 2, exit mileages spread uniformly across the boundary, and 52.3% of coincident units crossed into the penalized bracket.

Table 4. Decomposition of the coupling premium on coincident vehicles

Value component	Operational mechanism	Dollar value per coincident vehicle	Share of coupled advantage
Mileage steering $\Delta\pi_{\text{steer}}$	Upgrades, rate steering and dynamic dispatch	\$138.40	60.7%
Decision-focused loss $\Delta\pi_{\text{dl}}$	Alignment of the residual forecast with the disposal ranking	\$67.60	29.6%
Interaction $\Delta\pi_{\text{int}}$	Sensor-based condition tuning near the boundary	\$22.00	9.6%
Total net premium $\Delta\pi_{\text{coupled}}$	Coupled advantage, Arm 4 over Arm 2	\$228.00	100.0%

Decision-focused forecast training contributes \$67.60 per vehicle, 29.6% of the premium, and the interaction term adds \$22.00, 9.6%. Once a step-aware forecaster replaces the smooth one in sequential scoring, the decision-focused component falls to \$8.20 ($p = 0.34$) and the steering component stays at \$136.90. A forecaster that has learned the price drops reproduces the decision-focused component; steering needs control over dispatch.

4. Discussion

The evaluation separates two claims that the decision-learning literature holds in tension. Elmachoub and Grigas (2022) [13] trained prediction models on decision loss and reported improvements over predict-then-optimize on shortest-path and portfolio problems, particularly when the prediction model was misspecified. Hu et al. (2022) [16] derived fast rates for contextual linear optimization and showed that estimate-then-optimize attains them when the fitted model matches the data-generating relationship. A residual-value forecast that is smooth in mileage is misspecified at every price mark, and the estimated jumps of \$174.20 to \$218.60 give that misspecification a size in dollars. The profit results place its consequences where the theory puts them: coupling adds \$206 to \$228 per vehicle inside the coincidence band and \$7 to \$13 outside it, where neither difference departs from zero.

Two channels carry the premium, and they respond differently to forecast quality. Mileage steering supplies 60.7% of it. A sequential policy can change only the date of sale, and its exit mileages spread uniformly across the mark; a coupled policy also changes the rate at which mileage accrues, because pricing and assignment decide which reservations a vehicle serves, the lever that Hartman (2004) [3] describes as workload allocation. The decision-focused component supplies 29.6% and vanishes with a step-aware forecaster, whose gain over the smooth one is \$8.20 ($p = 0.34$). The steering gain stays at \$136.90 in that specification. Knowing the price at 80,000 miles does not put the odometer there.

Lessmann and Voß (2017) [17] rank resale-price forecasters by forecast accuracy, find random forest regression particularly effective, and advise against linear regression. Contrary to carrying that ranking into disposal decisions, the evidence indicates that forecasters should be ranked by decision regret near the marks. The smooth forecaster here has mean absolute out-of-sample errors of \$248.50 to \$281.00 at the three marks, and that error costs \$67.60 per coincident vehicle in decision regret; outside the band coupling changes profit by only \$7 to \$13. One accuracy figure averages the mileage range where the ranking is at stake with the range where it is not. The decision loss of Elmachoub and Grigas (2022) [13] scores the forecast where the ranking is made, and Dress et al. (2018) [18] add that forecast errors of opposite sign cost a lessor different amounts.

Sensor data and architecture contribute separately. Enriched telematics raise profit by \$86 to \$118 per vehicle in every cell of Table 3, which fits fleet-monitoring evidence that logged operating data reveal faults (Prytz et al., 2015; Rögnvaldsson et al., 2018; Theissler et al., 2021) [9,14,15]. Inside the band, coupling with odometer-only data earns \$3,618 per vehicle, \$110 more than sequential scoring with the full sensor stream at \$3,508. Outside the band the order reverses: the sensor stream adds \$86 and coupling adds \$7. A deployment that ships a lifecycle architecture together with sensors and reports one aggregate gain credits the architecture with income that the sensors earn on every vehicle. The shared pricing module (Besbes & Zeevi, 2009) [12] removes a third source from the comparison.

The two Kolesnykov articles map onto the arms as their designs are described. Revenue scoring on smoothed residual values (Kolesnykov, 2026b) [7] gives up \$206 per coincident vehicle to the coupled policy under odometer-only data (Kolesnykov, 2026a) [10] and gives up nothing measurable elsewhere, so the coupled framework earns its advantage on 28.2% of retirements. The interaction term of \$22.00 depends on the calibrated failure probabilities of the maintenance layer (Kolesnykov, 2026c) [8]. Predictive uncertainty degrades under dataset shift (Ovadia et al., 2019) [19], and concept drift in fleet mix or wholesale demand invalidates probabilities fitted earlier (Lu et al., 2019) [20], so the 2021 to 2025 estimate holds for as long as recalibration keeps pace with the fleet.

Four limits bound these estimates. The sample covers one operator, one vehicle class, the 65,000 to 95,000 mile range and three major marks. The operator sells through wholesale auctions, the setting of the price evidence (Lacetera et al., 2012) [1]; a buyback or dealer channel that prices mileage smoothly would make the smooth forecaster correctly specified, reduce e , and by equation (3) narrow the band. The half-width is a first-order approximation and holds where the retention margin is close to linear around $x_{d,i}$. The age dimension of the replacement threshold (Hartman, 2004) [3] carries a price discontinuity only if age-based marks exist in wholesale prices, and the evidence here documents odometer marks only. Steering also draws on reservation supply, such as short airport-transfer bookings, over the final 1,500 miles.

5. Conclusion

Coupling now has a measured price and a measured domain. At the operator studied, 28.2% of retirements sit inside the coincidence band, coupling adds \$206 to \$228 per vehicle there, and a step-aware forecaster reproduces only the \$67.60 decision-focused part of that premium. A fleet manager can compute w for each vehicle from three inputs, the forecast error, the margin slope and the distance to the next mark, and route the coupled policy to the vehicles that pass the test.

Whether steering survives a sales channel that prices mileage smoothly remains open. The proposition predicts a steering premium near zero there, because no jump exists to steer around. A test needs an operator that retires vehicles through buyback or dealer contracts, the same odometer-range restriction and matching procedure, equation (1) estimated on that channel to obtain γ_k , and Arms 1 and 3 run on matched cohorts. The proposition fails if a steering premium of the size found here appears where the estimated jumps are indistinguishable from zero.

References

- [1] Lacetera, N., Pope, D. G., & Sydnor, J. R. (2012). Heuristic thinking and limited attention in the car market. *American Economic Review*, 102(5), 2206-2236. <https://doi.org/10.1257/aer.102.5.2206>
- [2] Busse, M. R., Lacetera, N., Pope, D. G., Silva-Risso, J., & Sydnor, J. R. (2013). Estimating the effect of salience in wholesale and retail car markets. *American Economic Review*, 103(3), 575-579. <https://doi.org/10.1257/aer.103.3.575>
- [3] Hartman, J. C. (2004). Multiple asset replacement analysis under variable utilization and stochastic demand. *European Journal of Operational Research*, 159(1), 145-165. [https://doi.org/10.1016/S0377-2217\(03\)00397-7](https://doi.org/10.1016/S0377-2217(03)00397-7)

- [4] Oliveira, B. B., Carravilla, M. A., & Oliveira, J. F. (2017). Fleet and revenue management in car rental companies: A literature review and an integrated conceptual framework. *Omega*, 71, 11-26. <https://doi.org/10.1016/j.omega.2016.08.011>
- [5] Steinhart, C., & Gönsch, J. (2012). Integrated revenue management approaches for capacity control with planned upgrades. *European Journal of Operational Research*, 223(2), 380-391.
- [6] Oliveira, B. B., Carravilla, M. A., Oliveira, J. F., & Costa, A. M. (2019). A co-evolutionary matheuristic for the car rental capacity-pricing stochastic problem. *European Journal of Operational Research*, 276(2), 637-655.
- [7] Kolesnykov, V. (2026b). Revenue score modeling for optimal retention and disposal timing in rental vehicle fleets. *European Journal of Science, Innovation and Technology*, 6(4), 55-63.
- [8] Kolesnykov, V. (2026c). *Smart fleet management: Data-driven control of rental vehicles*. GlobeEdit.
- [9] Kolesnykov, V. (2026a). Demand-coupled fleet lifecycle optimization: A joint framework for pricing, maintenance scheduling, and disposition timing. *Austrian Journal of Technical and Natural Sciences*, (5-6), 283-290.
- [10] de Jonge, B., & Jakobsons, E. (2018). Optimizing block-based maintenance under random machine usage. *European Journal of Operational Research*, 265(2), 703-709. <https://doi.org/10.1016/j.ejor.2017.07.051>
- [11] Besbes, O., & Zeevi, A. (2009). Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57(6), 1407-1420. <https://doi.org/10.1287/opre.1080.0640>
- [12] Elmachoub, A. N., & Grigas, P. (2022). Smart "predict, then optimize." *Management Science*, 68(1), 9-26. <https://doi.org/10.1287/mnsc.2020.3922>
- [13] Hu, Y., Kallus, N., & Mao, X. (2022). Fast rates for contextual linear optimization. *Management Science*, 68(6), 4236-4245.
- [14] Lessmann, S., & Voß, S. (2017). Car resale price forecasting: The impact of regression method, private information, and heterogeneity on forecast accuracy. *International Journal of Forecasting*, 33(4), 864-877. <https://doi.org/10.1016/j.ijforecast.2017.04.003>
- [15] Dress, K., Lessmann, S., & von Mettenheim, H.-J. (2018). Residual value forecasting using asymmetric cost functions. *International Journal of Forecasting*, 34(4), 551-565.
- [16] Prytz, R., Nowaczyk, S., Rögnvaldsson, T., & Byttner, S. (2015). Predicting the need for vehicle compressor repairs using maintenance records and logged vehicle data. *Engineering Applications of Artificial Intelligence*, 41, 139-150.
- [17] Rögnvaldsson, T., Nowaczyk, S., Byttner, S., Prytz, R., & Svensson, M. (2018). Self-monitoring for maintenance of vehicle fleets. *Data Mining and Knowledge Discovery*, 32(2), 344-384. <https://doi.org/10.1007/s10618-017-0538-6>
- [18] Theissler, A., Pérez-Velázquez, J., Kettelgerdes, M., & Elger, G. (2021). Predictive maintenance enabled by machine learning: Use cases and challenges in the automotive industry. *Reliability Engineering & System Safety*, 215, Article 107864. <https://doi.org/10.1016/j.ress.2021.107864>
- [19] Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D., Nowozin, S., Dillon, J. V., Lakshminarayanan, B., & Snoek, J. (2019). Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift. *Advances in Neural Information Processing Systems*, 32.
- [20] Lu, J., Liu, A., Dong, F., Gu, F., Gama, J., & Zhang, G. (2019). Learning under concept drift: A review. *IEEE Transactions on Knowledge and Data Engineering*, 31(12), 2346-2363. <https://doi.org/10.1109/TKDE.2018.2876857>