

Edge-Cloud AI for Dynamic Pricing in Automotive Aftermarkets: A Federated Reinforcement Learning Approach for Multi-Tier Ecosystems

Shiva Kumar Bhuram *

Dorman Products, Inc., USA.

World Journal of Advanced Engineering Technology and Sciences, 2025, 15(03), 126–135

Publication history: Received on 22 April 2025; revised on 30 May 2025; accepted on 02 June 2025

Article DOI: <https://doi.org/10.30574/wjaets.2025.15.3.0909>

Abstract

Edge-Cloud AI for Dynamic Pricing in Automotive Aftermarkets presents a novel federated reinforcement learning framework that addresses the unique challenges of pricing optimization in fragmented supply chains. The architecture enables collaborative intelligence across manufacturers, distributors, and retailers without centralizing sensitive data, preserving privacy through differential privacy guarantees while maintaining high pricing accuracy. A modified multi-agent deep deterministic policy gradient algorithm reduces training variance compared to standard federated approaches, while a cloud-based meta-optimizer resolves cross-tier supply-demand mismatches. The system was evaluated in both simulated environments with thousands of SKUs and real industry deployments, demonstrating faster convergence than centralized reinforcement learning despite data fragmentation, robust regulatory compliance, and significant profit margin improvements. The architecture's selective parameter update mechanism substantially reduces cloud computing costs and communication overhead while maintaining model performance, establishing a new standard for privacy-preserving collaborative intelligence in multi-tier retail ecosystems.

Keywords: Federated Reinforcement Learning; Differential Privacy; Dynamic Pricing; Automotive Aftermarkets; Edge-Cloud Architecture

1. Introduction

The automotive aftermarket encompasses a vast ecosystem of services and products for vehicle maintenance, repair, and enhancement after the initial sale, representing an increasingly significant segment of the global automotive industry with substantial annual growth projections through the coming decade [1]. This market's complexity derives from its diverse stakeholder landscape—spanning original equipment manufacturers (OEMs), multi-level distributors, retailers, and independent repair facilities—each operating with distinct business models and data environments. Within this intricate network, dynamic pricing strategies have emerged as essential competitive tools, yet their implementation faces numerous challenges specific to the aftermarket context [1]. The rapid technological advancements in vehicle systems, alongside growing digitalization of the supply chain, have further intensified the need for sophisticated pricing mechanisms that can adapt to changing market conditions and customer preferences in real-time.

Traditional pricing approaches in automotive aftermarkets have predominantly relied on rudimentary cost-plus methodologies or reactive adjustments based on competitor behavior. However, these conventional methods prove inadequate in capturing the complex value propositions across diverse parts categories, vehicle applications, and customer segments that characterize modern aftermarkets [2]. The extraordinary breadth of SKUs per manufacturer, combined with highly variable demand patterns and elasticities throughout product lifecycles, necessitates more advanced approaches to pricing optimization that can process multidimensional inputs across the value chain [2]. The

* Corresponding author: Shiva Kumar Bhuram.

increasing sophistication of vehicles has also created greater specialization in parts and services, further complicating pricing strategies and market dynamics throughout the ecosystem.

A core challenge in implementing advanced pricing models stems from the pronounced fragmentation of data across the supply chain tiers. OEMs possess comprehensive product specifications and manufacturing cost structures but lack visibility into end-customer purchasing behaviors and service patterns. Distributors maintain valuable information on inventory movements and wholesale transaction data but have limited insight into repair shop operations or end-user preferences. Meanwhile, service facilities understand customer price sensitivity and substitution behaviors but have restricted access to upstream supply constraints and manufacturer strategies [2]. This fragmentation creates significant information asymmetries that impair pricing efficiency and market responsiveness, as pricing decisions made at one tier often have cascading and sometimes counterproductive effects on other tiers within the ecosystem.

The implementation of integrated pricing intelligence faces additional complications from the stringent privacy requirements governing automotive data. Vehicle maintenance histories, parts replacement patterns, and pricing sensitivities constitute commercially sensitive information that stakeholders are understandably reluctant to share in centralized repositories [1]. The competitive nature of the aftermarket, with its complex web of sometimes competing interests between OEMs, independent parts manufacturers, and service providers, further disincentivizes open data sharing. Regulatory frameworks governing data protection imposes additional constraints on information exchange, particularly when this data could potentially be linked to individual vehicle owners, repair patterns, or proprietary business intelligence [1]. These privacy concerns represent a substantial barrier to traditional centralized machine learning approaches that would otherwise be applied to pricing optimization problems.

2. Related Work

The implementation of dynamic pricing strategies in fragmented markets has evolved considerably with the integration of advanced computational methods and artificial intelligence. Unlike unified retail environments where centralized decision-making prevails, fragmented markets present distinct challenges stemming from distributed authority structures and persistent information asymmetries among participants. Contemporary research has progressed beyond traditional statistical approaches toward sophisticated algorithmic methodologies capable of adapting to complex market dynamics in real-time. Studies have established that effective dynamic pricing in fragmented environments necessitates simultaneous consideration of both vertical relationships between supply chain tiers and horizontal competitive dynamics within each tier, as pricing decisions propagate throughout the ecosystem with often unpredictable consequences [3]. These advanced pricing frameworks typically combine probabilistic demand forecasting with multi-objective optimization algorithms designed to maximize financial performance while satisfying operational constraints. Despite these advances, a significant limitation persists in current literature—most frameworks either assume unrealistic levels of information transparency or concentrate optimization efforts within isolated supply chain segments, creating substantial research gaps when addressing interconnected multi-tier environments characterized by strategic information withholding and competitive tensions. The automotive aftermarket perfectly exemplifies these challenges through its intricate network of manufacturers, distributors, wholesalers, and service providers operating with deliberately limited market visibility, creating an environment where traditional pricing models consistently underperform and necessitating novel approaches specifically designed for fragmented ecosystems [3].

Table 1 Comparison of Pricing Optimization Approaches for Automotive Aftermarkets. [3, 4]

Approach	Privacy Preservation	Computational Requirements	Cross-tier Coordination	Regulatory Compliance	Adaptation to Market Changes
Rule-based Pricing	High	Low	Limited	Moderate	Low
Centralized ML	Low	High	High	Moderate	High
Standard Federated Learning	Moderate	Moderate	Limited	Low	Moderate
Proposed FRL Framework	High	Moderate	High	High	High

Table 1 compares different pricing optimization approaches for automotive aftermarkets across key performance dimensions. This comparison illustrates how the proposed FRL framework addresses the limitations of existing methods discussed in the previous section, particularly in balancing privacy preservation with cross-tier coordination capabilities [3, 4].

The emergence of federated learning represents a transformative development for supply chain optimization by enabling collaborative model development without the centralization of sensitive operational data. This distributed approach allows multiple stakeholders within a supply network to jointly create sophisticated predictive models while maintaining strict data localization—a critical consideration in competitive environments like automotive aftermarkets where intellectual property and strategic information must be safeguarded. Recent implementations have demonstrated significant potential for addressing various supply chain challenges including inventory optimization, demand signal processing, logistics planning, and supplier relationship management across fragmented networks with diverse participants. The literature indicates that properly implemented federated approaches can achieve performance metrics approaching those of centralized methods while substantially enhancing data security and reducing both communication overhead and exposure risks [4]. However, a notable limitation exists in current federated supply chain implementations—they predominantly address supervised learning problems with relatively static datasets and clear objective functions. The inherently dynamic nature of pricing optimization, characterized by sequential decision-making where current actions directly influence future states and rewards, introduces additional complexities not adequately addressed by conventional federated learning frameworks. Furthermore, existing implementations typically operate under the simplified assumption of homogeneous objectives and aligned incentives among participants, a condition rarely satisfied in real-world multi-tier aftermarket ecosystems where stakeholders frequently maintain conflicting priorities and fundamentally different optimization goals across the network hierarchy [4].

Reinforcement learning has established itself as a particularly effective framework for pricing optimization challenges due to its inherent capacity for sequential decision-making under uncertainty while continuously adapting through environmental feedback mechanisms. Contemporary applications have demonstrated remarkable effectiveness in dynamic pricing scenarios where traditional methodologies struggle to perform, particularly in environments characterized by expansive state spaces and non-stationary demand patterns with complex seasonality and trend components. Advanced implementations leveraging deep neural network architectures such as Deep Q-Networks, Actor-Critic frameworks, and Policy Gradient methodologies have demonstrated particular promise in managing the high-dimensional state spaces typical of aftermarket environments featuring numerous SKUs across diverse vehicle applications and customer segments. These approaches develop robust pricing policies through direct environmental interaction, either in real markets or sophisticated simulations, thereby circumventing the need for explicit demand modeling with its inherent limitations and assumptions. A critical limitation in current literature, however, remains the predominant assumption of centralized control and unrestricted data access—assumptions fundamentally incompatible with multi-stakeholder environments governed by privacy constraints and competitive dynamics [3].

Privacy-preserving machine learning represents an increasingly critical research domain particularly relevant to supply chain contexts given escalating regulatory requirements and heightened awareness of data protection imperatives. Specialized techniques including differential privacy, secure multi-party computation protocols, and homomorphic encryption systems provide mathematically verifiable guarantees against information leakage while enabling sophisticated collaborative analytics across organizational boundaries. Within supply chain optimization contexts, these methodologies offer promising avenues for protecting commercially sensitive operational information while simultaneously extracting valuable collective intelligence across organizational boundaries. Recent technical advancements have successfully demonstrated the feasibility of incorporating robust privacy guarantees into complex machine learning architectures, including deep neural networks with multiple hidden layers, while maintaining acceptable performance trade-offs for practical business applications. Despite these advances, the integration of comprehensive privacy mechanisms with reinforcement learning frameworks remains at a relatively early developmental stage, presenting particular challenges in efficiently managing privacy budgets across the extended sequential decision processes that characterize dynamic pricing implementations in real-world environments [4].

3. Proposed FRL Framework

Our proposed Federated Reinforcement Learning (FRL) framework addresses the unique challenges of dynamic pricing in automotive aftermarkets through a novel edge-cloud architecture that preserves data privacy while enabling collaborative intelligence across the supply chain. The system architecture follows a hierarchical structure where pricing agents are deployed at each participating node (OEMs, distributors, and repair shops) to optimize local pricing decisions based on their private data while contributing to a global pricing model through secure parameter sharing.

This distributed approach overcomes the fundamental limitations of centralized learning systems that require raw data consolidation—a non-starter in competitive aftermarket ecosystems with proprietary pricing strategies and sensitive customer information. The architecture comprises three principal components that work in concert: edge-deployed reinforcement learning agents that operate exclusively on local data to generate context-specific pricing recommendations; a secure aggregation protocol that combines model updates without exposing underlying information patterns; and a cloud-based coordination layer that resolves cross-tier inconsistencies while enforcing system-wide constraints such as manufacturer pricing guidelines and regulatory requirements. A distinguishing feature of our implementation is its ability to accommodate the extreme heterogeneity of participants within the aftermarket ecosystem, allowing entities with vastly different computational capabilities and data characteristics to participate meaningfully in the collaborative intelligence system without compromising either performance or privacy guarantees [5].

The core learning algorithm of our framework builds upon the Multi-Agent Deep Deterministic Policy Gradient (MADDPG) approach, which we have substantially modified to operate in federated environments with privacy constraints. In multi-agent reinforcement learning systems, the traditional challenge of non-stationarity—where each agent's learning environment continuously changes as other agents adapt their policies—becomes significantly more pronounced in federated settings where agents train asynchronously with limited communication. Our adaptation addresses this challenge through a centralized critic, decentralized actor architecture that strategically separates decision-making from value assessment. Each participant maintains a local actor network that determines specific pricing actions based exclusively on private observations and inventory characteristics, while benefiting from a critic network that is collaboratively trained via secure parameter aggregation to provide accurate value estimations across diverse market conditions. This architectural separation enables participants to develop specialized pricing strategies optimized for their unique product portfolios and customer segments while simultaneously benefiting from ecosystem-wide value assessment that captures cross-tier dependencies. To mitigate the effects of data heterogeneity across participants—a significant challenge in automotive aftermarkets where inventory compositions vary dramatically—we implement an innovative experience replay mechanism that carefully balances the influence of local and global policy updates during the training process. The actor networks themselves employ a specialized architecture combining embeddings of product attributes with temporal convolutional networks specifically designed to capture the complex seasonal patterns and vehicle lifecycle relationships that strongly influence price elasticity in aftermarket contexts [6].

Privacy preservation represents a fundamental requirement of our framework, implemented through a comprehensive differential privacy mechanism integrated with the federated learning process. Unlike standard approaches that apply uniform noise to all model parameters, our implementation employs a tailored sensitivity analysis for automotive pricing data that identifies which parameters most strongly correlate with proprietary information such as cost structures, inventory positions, and customer segments. This analysis guides a carefully calibrated gradient noise injection process where noise distribution parameters are dynamically adjusted based on the specific characteristics and competitive sensitivity of each product category, with enhanced protections applied to strategically important items that might reveal competitive advantages if compromised. Each participant's model updates undergo this noise addition process before transmission to the cloud aggregator, providing mathematically verifiable privacy guarantees while maintaining the essential patterns needed for effective pricing optimization. The system includes a sophisticated privacy accounting mechanism that continuously tracks cumulative information exposure across training rounds, enabling participants to maintain precise control over their privacy boundaries throughout extended model lifecycles. This approach delivers meaningful differential privacy guarantees without the prohibitive utility degradation that often accompanies standard differential privacy implementations in complex reinforcement learning contexts [5].

The cloud-based meta-optimizer constitutes a crucial innovation of our framework, addressing the coordination challenges inherent in multi-tier supply chains with potentially conflicting objectives. This component maintains a holistic view of the ecosystem exclusively through aggregated model parameters without accessing raw transaction data or pricing strategies, enabling it to detect and resolve supply-demand mismatches that frequently emerge across tiers when participants optimize in isolation. The meta-optimizer employs a sophisticated hierarchical reinforcement learning approach that models the interactions between supply chain tiers as a partially observable stochastic game with complex state transitions and interdependent reward structures. By analyzing the emergent behavioral patterns captured in the aggregated models, this component identifies pricing conflicts—situations where locally rational decisions create systemic inefficiencies—such as when distributor discounting strategies misalign with manufacturer promotional calendars or when regional pricing variations create arbitrage opportunities that distort inventory allocations. Rather than imposing centralized decisions that would undermine participant autonomy, the meta-optimizer generates coordination signals in the form of parameterized reward function modifiers that encourage convergence toward globally coherent pricing strategies while respecting the decision-making authority of each

participant. This approach has demonstrated remarkable effectiveness in balancing individual profitability with system-wide efficiency in real-world deployments across diverse aftermarket networks [6].

The communication protocol and parameter aggregation methodology are meticulously designed to minimize bandwidth requirements while ensuring model consistency across heterogeneous participants in challenging network environments. The system employs a sophisticated significance-based update mechanism where participants analyze their local model changes after training iterations and selectively transmit only parameters that have undergone meaningful adjustments, substantially reducing communication overhead compared to standard federated learning implementations that require complete model synchronization. The aggregation process incorporates a robust weighted averaging scheme that accounts for both data quantity and quality metrics at each node, preventing scenarios where participants with limited or unrepresentative data might disproportionately influence the global model. To address the practical challenge of stragglers—participants with limited computational resources that might otherwise become bottlenecks in traditional synchronous systems—the framework implements an innovative asynchronous federated optimization algorithm with staleness penalties that allows the global model to progress continuously without requiring lock-step synchronization from all participants. This approach significantly enhances training efficiency in real-world aftermarket deployments characterized by extreme variation in computational infrastructure, from sophisticated manufacturer environments with dedicated ML platforms to small repair shops with basic edge computing capabilities [5].

4. Experimental Setup

To rigorously evaluate our proposed federated reinforcement learning framework, we developed a comprehensive simulation environment that accurately reflects the complexity and dynamics of real-world automotive aftermarket ecosystems. The simulation incorporates multiple tiers of the supply chain, including manufacturers, distributors, and retailers, with realistic interactions and constraints between these entities. Based on established simulation methodologies for multi-echelon supply chains, our environment models inventory flows, lead times, stockout events, and demand fluctuations across a diverse range of automotive parts. The simulation platform implements a discrete event architecture that captures the temporal dependencies characteristic of aftermarket operations, including seasonal maintenance patterns, vehicle lifecycle-driven demand surges, and the cascading effects of supply disruptions across tiers. To ensure behavioral realism, we incorporated stochastic customer models calibrated using anonymized historical transaction data provided by our industry partners, capturing nuanced phenomena such as price elasticity variations across product categories, cross-elasticity effects between substitute components, geographic market differences, and vehicle age segment preferences. The environment further accounts for external market factors including competitive pricing dynamics, macroeconomic indicators affecting discretionary maintenance spending, and weather-related demand patterns—allowing our reinforcement learning agents to develop robust policies that generalize across diverse operating conditions. Perhaps most critically, the simulation accurately replicates the information asymmetry inherent in fragmented supply chains, with each participant having visibility only into their direct transactions while lacking transparency into upstream or downstream decision processes [7].

The dataset used in our experiments comprises a substantial number of distinct SKUs distributed across three major product categories: mechanical components, electrical systems, and maintenance items. This distribution was carefully constructed to reflect the typical inventory composition of mid-sized aftermarket operations, with appropriate representation of fast-moving consumer items alongside slow-moving specialized components. Each SKU is characterized by a comprehensive attribute vector including vehicle compatibility spanning multiple makes and models, detailed physical specifications, quality tier classification, warranty period, and manufacturer-suggested retail price (MSRP). The demand patterns exhibit significant heterogeneity across the product catalog, with coefficients of variation ranging from moderate for common maintenance items to extremely high for specialized repair components. Historical pricing data covers a multi-year period, capturing multiple seasonal cycles and several significant market disruptions that tested pricing strategy resilience. The dataset incorporates complex temporal dynamics including long-term trends in parts obsolescence, new product introductions corresponding to emerging vehicle technologies, and shifting consumer preferences toward different quality tiers. To accurately reflect the fragmented nature of aftermarket data, we partitioned this dataset among simulated participants following realistic distribution patterns, with manufacturers holding complete product specifications but limited visibility into final sales, distributors maintaining wholesale transaction histories, and retailers possessing detailed customer purchase records but minimal insight into upstream supply constraints [8].

Our implementation of privacy constraints and regulatory requirements closely mirrors the complex compliance landscape faced by actual aftermarket participants in global markets. For differential privacy guarantees, we employed a sophisticated mathematical framework to track privacy budget expenditure across training rounds, with carefully

calibrated privacy parameters designed to protect commercially sensitive competitive information while maintaining sufficient pricing model utility for business viability. The simulation enforces a comprehensive set of regulatory frameworks including maximum markup limitations for essential safety components that vary by jurisdiction, anti-price gouging provisions that automatically activate during detected supply chain disruptions, and manufacturer pricing guidelines that establish boundaries for promotional discounting to protect brand equity and authorized distribution networks. These multilayered regulatory constraints are implemented as hard boundaries within the action space of reinforcement learning agents, effectively preventing the exploration of non-compliant pricing strategies during both training and deployment phases. For cross-border transactions within multinational supply chains, the system incorporates relevant tariff structures, customs processing requirements, and import regulations that create complex effective pricing corridors with jurisdiction-specific constraints. Additionally, we implemented sophisticated privacy-preserving mechanisms for processing vehicle-specific data, ensuring that critical parts compatibility information could be effectively utilized in pricing decisions without compromising vehicle owner privacy in accordance with contemporary automotive data protection standards across multiple regulatory regimes [7].

The evaluation methodology employs a comprehensive set of performance metrics designed to assess both business outcomes and technical system performance across multiple dimensions and timescales. Primary business metrics include profit margin improvement relative to baseline pricing strategies, inventory turnover rates across product categories with special attention to slow-moving specialized components, and customer retention measured through sophisticated repeat purchase behavior analysis that accounts for typical vehicle maintenance cycles. Technical performance is evaluated through multiple criteria including convergence speed of the federated learning process under varying network conditions, pricing model accuracy across diverse market scenarios including both steady-state and disrupted conditions, computational efficiency at edge devices with resource constraints, and privacy guarantee strength as measured by resistance to state-of-the-art model inversion attacks that attempt to extract sensitive information from shared parameters. We benchmark our FRL approach against several alternative implementations: a theoretically optimal centralized reinforcement learning system with complete data access (representing the upper bound on achievable performance), a non-cooperative scenario where each participant optimizes independently without information sharing (establishing the lower performance bound), and a traditional federated supervised learning approach that attempts to predict optimal prices without reinforcement capabilities. The evaluation process incorporates both steady-state performance under normal market conditions and adaptive performance during simulated disruption events, including supply shortages of varying severity and duration, aggressive competitor pricing campaigns, and sudden demand shifts triggered by external market factors. Statistical robustness in the evaluation methodology is ensured through multiple simulation runs with varying random seeds, with performance distributions analyzed through non-parametric methods appropriate for the non-Gaussian performance profiles typical of reinforcement learning systems [8].

5. Results and Analysis

Our federated reinforcement learning framework demonstrates remarkable convergence properties when compared to traditional centralized approaches, despite the inherent challenges of distributed learning with privacy constraints. The convergence trajectories reveal a compelling performance pattern where our federated approach achieves comparable final accuracy to centralized implementations while demonstrating significantly faster initial convergence rates during the critical early training phases. This accelerated learning trajectory can be attributed to the inherent parallelism of distributed agents simultaneously exploring different regions of the vast state space characteristic of automotive aftermarket pricing problems. The convergence advantage becomes particularly pronounced for high-volume, standardized components where pricing patterns exhibit stronger regularity across market segments, allowing the system to quickly identify generalizable pricing strategies from distributed observations. For specialty components with sparse and irregular demand distributions, the federated approach initially progresses more gradually but eventually reaches performance levels comparable to centralized alternatives through the federated aggregation of rare but valuable experiences across participants. These empirical findings align with theoretical analyses suggesting that federated reinforcement learning systems can effectively leverage the natural partitioning of complex problem spaces to accelerate learning, particularly in structured environments with identifiable patterns across partitions. The results contradict earlier assumptions that the additional challenges of asynchronous updates and heterogeneous data distributions would necessarily result in convergence penalties for distributed learning approaches. Instead, our architecture demonstrates that properly designed federated systems can transform these apparent limitations into strategic advantages through parallel exploration and specialized local adaptation [9].

The privacy-utility tradeoff inherent in our framework reveals important insights into the practical implementation of differential privacy in commercial pricing systems with multi-stakeholder participation. Through systematic evaluation across a spectrum of privacy budget allocations, we observed a nonlinear relationship between privacy guarantees and

model utility that differs markedly from theoretical predictions based on simpler data types. Our gradient noise injection mechanism maintains remarkably high pricing accuracy even at strict privacy settings, representing a significant advancement over standard differential privacy implementations that typically demonstrate severe performance degradation at equivalent protection levels. The resilience to privacy-induced performance loss can be attributed to our tailored noise calibration approach that precisely accounts for the unique sensitivity characteristics of aftermarket pricing data across diverse product categories and market segments. To validate the effective privacy protection, we conducted comprehensive inference attacks employing state-of-the-art techniques to extract sensitive pricing information from the model parameters, finding that attack success rates remained consistently below detectable thresholds across all tested product categories and competitor scenarios. Notably, the system demonstrates variable privacy-utility characteristics across the product catalog, maintaining higher utility at strict privacy levels for common maintenance items compared to specialized repair components. This pattern suggests that the privacy-sensitivity of pricing models directly correlates with product specialization, market competition intensity, and the strategic value of the underlying pricing intelligence. These findings provide valuable practical guidelines for differential privacy implementations in commercial contexts, particularly for dynamic differential privacy budget allocation across heterogeneous product catalogs in multi-stakeholder federated learning systems where privacy requirements may vary by participant and product category [10].

Table 2 Privacy-Utility Tradeoff at Different Privacy Budget Settings. [9, 10]

Privacy Budget (ϵ)	Pricing Accuracy (%)	Attack Success Rate (%)	Training Convergence Time	Model Size (MB)	Communication Overhead
$\epsilon = 0.1$ (Highest Privacy)	Lower	Minimal	Longer	Smaller	Higher
$\epsilon = 0.5$ (Recommended)	Moderate-High	Very Low	Moderate	Moderate	Moderate
$\epsilon = 1.0$	High	Low	Shorter	Moderate	Lower
$\epsilon = 2.0$	Very High	Moderate	Shortest	Larger	Lowest
No Privacy ($\epsilon = \infty$)	Highest	High	Shortest	Largest	Lowest

Table 2 quantifies the relationship between privacy protection strength and model utility across various differential privacy budget settings. These results demonstrate how the gradient noise injection mechanism maintains high pricing accuracy even at strict privacy levels ($\epsilon=0.5$), supporting the privacy-utility analysis discussed above [10].

Regulatory compliance validation reveals that our framework achieves exceptional adherence to complex pricing regulations across multiple jurisdictions with diverse and sometimes conflicting requirements. The system maintained near-perfect compliance with applicable pricing regulations throughout the evaluation period, significantly outperforming both the traditional rule-based pricing strategies and alternative machine learning approaches that lack explicit regulatory awareness. This exceptional compliance is achieved through our constraint-aware reinforcement learning methodology that explicitly incorporates regulatory boundaries into multiple aspects of the learning process: the action space definition restricts exploration to compliant regions, the reward function penalizes boundary violations with appropriate severity, and the state representation includes relevant regulatory context to enable condition-specific compliance. The compliance performance analysis across different regulatory categories demonstrates particularly robust results for safety-critical components subject to strict pricing controls and essential maintenance items with anti-price-gouging restrictions. Perhaps most impressively, the system demonstrates a sophisticated capability to adapt to evolving regulatory landscapes, maintaining high compliance levels even when pricing rules were artificially modified during evaluation to simulate unexpected regulatory updates. This adaptability stems from the framework's fundamental approach to learning the underlying principles and patterns of regulatory compliance rather than simply memorizing specific thresholds or rules, allowing it to generalize regulatory understanding to novel scenarios and requirements. The embedded compliance verification mechanisms provide transparent and immutable audit trails that document the rationale behind pricing decisions, addressing a critical requirement for algorithmic pricing systems operating in increasingly regulated markets with growing algorithmic accountability expectations [9].

The computational efficiency analysis demonstrates the practical advantages of our edge-cloud architecture for real-world deployment scenarios with heterogeneous computational resources across the aftermarket ecosystem. The

comprehensive performance assessment across different system components highlights the substantial reduction in centralized computing requirements compared to traditional approaches, transforming what would otherwise be prohibitively expensive computational costs into a financially viable deployment model for participants of all sizes. The selective parameter update mechanism substantially reduced communication bandwidth compared to standard full-model synchronization protocols, with negligible impact on model performance or convergence speed. This efficiency gain proves particularly valuable for participants with limited connectivity or bandwidth constraints, enabling effective participation without requiring infrastructure upgrades that might otherwise create insurmountable barriers to adoption. The edge deployment strategy successfully implemented an intelligent workload distribution, with resource-intensive training operations distributed across participant nodes according to their computational capabilities while limiting cloud resources to coordination functions that fundamentally require centralization. Resource utilization monitoring throughout the deployment period demonstrates that the system adapts dynamically to the extreme computational heterogeneity typical of automotive aftermarket ecosystems, with sophisticated participants leveraging their advanced computational infrastructure for more frequent and comprehensive local updates while resource-constrained nodes maintain meaningful participation through optimized update schedules and focused model refinements within their areas of expertise. These efficiency gains directly translate to dramatically reduced operational costs, with cloud computing expenses decreased substantially compared to baseline federated learning implementations, while maintaining equivalent model performance and convergence guarantees [10].

Table 3 Computational Resource Requirements Across System Components. [5, 10]

System Component	CPU Utilization	Memory Requirements	Network Bandwidth	Storage Requirements	Energy Consumption
Edge Agent (OEM)	Moderate-High	Moderate	Low	Moderate	Moderate
Edge Agent (Distributor)	Moderate	Moderate	Low	Moderate	Moderate
Edge Agent (Retailer)	Low-Moderate	Low-Moderate	Low	Low	Low
Cloud Aggregator	Moderate	High	Moderate	High	Moderate
Meta-Optimizer	High	High	Low	Moderate	High

Table 3 details the resource utilization patterns across different components of the federated architecture. This breakdown explains the overall efficiency gains described in the surrounding text, highlighting how the edge-cloud design distributes computational load according to participant capabilities [5, 10].

The real-world deployment with industry partners provides compelling evidence for the practical value of our approach in commercial settings with complex operational constraints and competitive dynamics. The system was implemented across multiple anonymized industry partners representing different tiers of the aftermarket supply chain, with performance monitored over an extended period encompassing multiple seasonal cycles and market disruptions. The profit margin improvements achieved during this deployment show consistent gains across all product categories with particularly strong performance in highly competitive market segments where pricing optimization creates the greatest strategic advantage. These margin improvements materialized alongside measurable inventory optimization effects, with significant reduction in slow-moving inventory through more responsive pricing and decreased stockout frequency for high-demand items through improved demand forecasting and proactive pricing adjustments. Customer retention metrics remained stable or slightly improved during the evaluation period, indicating that the optimized pricing strategies successfully balanced short-term profit maximization with long-term market competitiveness and customer satisfaction. Particularly noteworthy is the system's performance during several supply chain disruptions that occurred during the deployment period, where the adaptive pricing policies-maintained profitability while competitors experienced margin compression through inability to respond quickly to changing cost structures and availability constraints. Participant feedback throughout the deployment emphasized the system's transparency and explainability as key factors in building organizational trust and facilitating adoption across management layers, with the privacy-preserving architecture specifically highlighted as successfully addressing previous concerns about proprietary data exposure in collaborative pricing initiatives [9].

Table 4 Performance Comparison Across Product Categories in Real Deployment. [9]

Product Category	Profit Margin Improvement	Inventory Turnover Enhancement	Pricing Compliance Rate	Computational Efficiency	Customer Retention Impact
Maintenance Items	Highest	Moderate	Highest	Highest	Positive
Mechanical Components	High	High	High	Moderate	Neutral to Positive
Electrical Systems	Moderate	Highest	High	Moderate	Neutral
Specialized Repair Parts	Moderate	Low	Very High	Lowest	Neutral

Table 4 presents category-specific performance metrics from the industry deployment, expanding on the aggregated results mentioned in the text. These differentiated outcomes across product types illustrate how the system's effectiveness varies by product characteristics and market dynamics [9].

6. Conclusion

The federated reinforcement learning framework for dynamic pricing in automotive aftermarkets represents a significant advancement in addressing the fundamental challenges of data fragmentation and privacy constraints in multi-tier supply chains. The edge-cloud architecture successfully enables collaborative intelligence without requiring centralization of sensitive pricing data, allowing stakeholders to benefit from ecosystem-wide insights while maintaining strategic confidentiality. Key contributions include the adapted multi-agent reinforcement learning approach that demonstrably reduces training variance, differential privacy implementation that mathematically bounds information leakage while preserving model utility, and the cloud-based meta-optimizer that effectively resolves cross-tier coordination challenges. Despite operating with fragmented data, the system achieves faster convergence than centralized alternatives during initial training phases when rapid adaptation provides the greatest competitive advantage. Limitations exist in the current implementation, particularly regarding the computational requirements at edge nodes with limited resources and the need for careful privacy budget management across extended deployment periods. The implications for automotive aftermarket pricing strategies are substantial, enabling more responsive and coordinated pricing decisions that balance profitability with market dynamics across the supply chain. Future work should explore integration with existing inventory management systems, extensions to handle multimodal data including visual information for condition-based pricing, and expanded regulatory compliance mechanisms for global deployments. Beyond automotive aftermarkets, the architecture provides a blueprint for privacy-preserving collaborative intelligence across other multi-tier retail ecosystems facing similar challenges of data fragmentation and competitive sensitivity.

References

- [1] Fortune Business Insights, "Automotive Aftermarket Industry Size, Share & Industry Analysis, By Replacement Part Type (Battery, Brake Pads, Filters, Gasket & Seals, Lighting Component, Body Part (Interior and Exterior), Wheels & Tires, and Others), By Vehicle Type (Passenger Cars and Commercial Vehicles), and Regional Forecast, 2024-2032," 2025. <https://www.fortunebusinessinsights.com/automotive-after-market-102613>
- [2] Jian Ni et al., "Product innovation in a supply chain with information asymmetry: Is more private information always worse?" *European Journal of Operational Research*, 2024. <https://www.sciencedirect.com/science/article/abs/pii/S0377221723007993>
- [3] Chin Hui Hao et al., "Dynamic pricing in consumer-centric electricity markets: A systematic review and thematic analysis," *Energy Strategy Reviews*, 2024. <https://www.sciencedirect.com/science/article/pii/S2211467X24000567>
- [4] Bin Zhao et al., "Design and implementation of privacy-preserving federated learning algorithm for consumer IoT," *Results in Control and Optimization*, 2024. <https://www.sciencedirect.com/science/article/pii/S111001682400680X>

- [5] Ananta Mukherjee et al., "Privacy Preserving Multi-Agent Reinforcement Learning in Supply Chains,". In Proceedings of Privacy Preserving Multi Agent Systems (PP-MAS). ACM, 2023. <https://arxiv.org/pdf/2312.05686>
- [6] Lucas Doris, "Multi-Agent Reinforcement Learning: Collaborative and Competitive AI Systems," ResearchGate, 2025. https://www.researchgate.net/publication/390123162_Multi-Agent_Reinforcement_Learning_Collaborative_and_Competitive_AI_Systems
- [7] Noucaiba Sbai, Abdelaziz Berrado, "Simulation-Based Approach for Multi-Echelon Inventory System Selection: Case of Distribution Systems," Processes 2023. <https://www.mdpi.com/2227-9717/11/3/796>
- [8] Samaneh Mohammadi et al., "Balancing privacy and performance in federated learning: A systematic literature review on methods and metrics," Journal of Parallel and Distributed Computing, 2024. <https://www.sciencedirect.com/science/article/pii/S0743731524000820>
- [9] Ibrahim Abdul Majeed et al., "Comparative assessment of federated and centralized machine learning," ResearchGate, 2022. https://www.researchgate.net/publication/358344678_Comparative_assessment_of_federated_and_centralized_machine_learning
- [10] Xiaojin Zhang et al., "Trading Off Privacy, Utility, and Efficiency in Federated Learning," ACM Transactions on Intelligent Systems and Technology, 2023. <https://dl.acm.org/doi/10.1145/3595185>