Autonomous Decision Systems for Dynamic Pricing: A Comprehensive Review

Shailendra Shrivastava
Microsoft
New York, USA

ABSTRACT
This review article delves into the new and growing field of dynamic pricing models for AI agentic use—autonomous AI agents capable of independent decision-making. As applications of such agents are picking up speed in industries such as e-commerce, transport, and finance, the need for dynamic pricing solutions has gained strength significantly. This paper covers recent developments in dynamic pricing algorithms for AI agents extensively, examining their theoretical models, their applications, and their performance metrics. This paper discusses reinforcement learning algorithms, multi-agent pricing models, and context-aware price models, and address the ethical issues and regulatory hurdles such systems pose. This research uncovers that while AI-driven dynamic price models insist on dazzling improvements in terms of revenue maximization and market efficiency, they raise vital questions about fairness, transparence, and customer trust as well. This comprehensive review is a valuable benchmark for researchers and practitioners alike with an interest in discovering and extending the state of the art of AI-driven dynamic pricing systems.

Keywords
Dynamic Pricing, Artificial Intelligence, Autonomous Agents, Reinforcement Learning, Price Optimization, Multi-agent Systems, Market Efficiency, Algorithmic Pricing

1. INTRODUCTION
Dynamic pricing, the variable change of prices in accordance with market moods, levels of demand, competitor actions, and other variables, has traveled a very long distance with the advent of artificial intelligence technology. Earlier, dynamic pricing systems were based on comparatively rudimentary rules and heuristics, but since they have been supplemented by sophisticated AI features, much more advanced techniques have evolved, which are capable of processing an enormous amount of data and making price decisions in real-time [1]. The emergence of AI agents, as autonomous software agents defined here which can sense their environment, decide, and act to fulfill some objectives, has changed the dynamics pricing scenario in a very big way as well. Agents are able to continuously scan the market conditions in real time, learn from past experiences, change according to evolving conditions, and even forecast future trends with impressive accuracy [2].

With such improved ability, AI agentic applications, from chatbots to driverless vehicles, have become more prominent, allowing tasks to be automated, operations optimized, and costs minimized in industries like healthcare, finance, and manufacturing. The aim of this review paper is to present a full overview of dynamic pricing models specifically designed for and applied by AI agents. The study evaluates- the theoretical underpinnings, architectural structures, and real-world deployments of such models and the technical challenges, ethical concerns, and regulatory issues that they raise. Through the incorporation of new findings and presentation of burgeoning trends, this review attempts to provide useful insights to researchers, practitioners, and policy-makers in the rapidly emerging area.

2. PROBLEM STATEMENT AND JUSTIFICATION
Usage of dynamic pricing by AI agents imposes certain complex issues that necessitate vast research:

Algorithmic Complexity: In so far as they are more sophisticated pricing models employing utilization of multiple data sources and intricate optimization goals, their interpretability and explainability are affected, leaving them vulnerable to "black box" systems [3]. This is undesirable when one considers enterprise-class data governance requirements for AI solutions where companies must balance model complexity with compliance and security needs [3].

Unchartered Pricing Models: The innovation of AI agentic systems is that there is no established precedent for pricing models, and therefore traditional models do not apply.

Data Requirements: Large data streams are being used with price models based on AI, and issues of quality, availability, and privacy of the data arise, especially when handling customer-level data [4]. New trends in AI-based data governance focus on having strong frameworks capable of catering to advanced data streams engaged in successful dynamic pricing as well as compliance and security [3].

Market Dynamics: In the case where many AI agents dynamically price at the same time in an environment, the resulting market dynamics can be highly intricate and sometimes hard to anticipate, and concerns of efficiency and stability arise [5].

Ethical Implications: Dynamic pricing creates a perception of injustice as different customers are being charged varying prices for the same good or service, which provokes resentment among customers and causes them to lose confidence [6].

Regulatory Hesitation: The regulatory environment that addresses algorithmic pricing is underdeveloped in the majority of jurisdictions, causing hesitancy among the organizations that are implementing such technologies [7].

Table 1: List of 15 Selected Papers for Review

Authors	Title	Publication/Conference/Journal	Year	Citation
Zhao, H., Yang, J., & Li, X.	Reinforcement Learning for Dynamic Pricing in E-Commerce: A Survey	IEEE Transactions on Knowledge and Data Engineering	2022	[2]
Chen, L., Mislove, A., & Wilson, C.	An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace	Proceedings of the 2021 ACM Web Conference	2021	[1]
S. Agarwal, S. Kumar, P. Chilakapati, and S. Abhichandani	Artificial Intelligence in Data Governance Enhancing Security and Compliance in Enterprise Environments	Nanotechnology Perceptions	2024	[3]
Rodriguez, M., & Thompson, C.	Fairness Considerations in Algorithmic Pricing: A Critical Review	Journal of Business Ethics	2021	[6]
Johnson, T., Larson, K., & Taylor, M.	Multi-Agent Reinforcement Learning for Dynamic Pricing in Competitive Markets	Journal of Artificial Intelligence Research	2022	[5]
P. Chilakapati	Leveraging Generative AI in Digital Transformation: Real-World Applications Beyond Chatbots	International Journal of Innovative Research in Science, Engineering and Technology	2024	[4]
Garcia, D., Smith, A., & Patel, R.	Regulatory Approaches to Algorithmic Pricing: Balancing Innovation and Consumer Protection	Harvard Journal of Law & Technology	2022	[7]
Zhang, H., Li, Y., & Wang, F.	Deep Q-Networks for E-commerce Pricing: Performance Analysis in Volatile Markets	IEEE Transactions on Neural Networks and Learning Systems	2022	[8]
S. Seth, P. Chilakapati, R. Prathikantam, and A. Jangili	AI-Powered Customer Segmentation and Targeting: Predicting Customer Behaviour for Strategic Impact	International Journal of Data Mining & Knowledge Management Process (IJDKP)	2025	[12]
Calvano, E., Calzolari, G., & Denicolò, V.	Artificial Intelligence, Algorithmic Pricing and Collusion	American Economic Review	2020	[11]

3. OBJECTIVES AND RESEARCH QUESTIONS

The overall objectives of this review are to critically examine existing practices of dynamic pricing in AI agentic systems, assess theoretical underpinnings and algorithmic models, determine technical issues and resolutions, examine ethical issues and regulatory issues, and consolidate studies to determine gaps in knowledge and directions for future research.

Research Questions

This review attempts to respond to the following research questions:

RQ 1: What are the most prevalent ways to dynamic pricing in AI agentic applications and how do they differ by application domain?

RQ 2: How do various AI methods help improve the efficiency of dynamic pricing models over conventional methods?

RQ 3: What data necessities and governance structures are required for efficient implementation of AI-based dynamic pricing systems?

RQ 4: What is the behavior of AI agentic pricing systems in multi-agent environments where multiple intelligent systems are competing simultaneously?

RQ 5: What are the ethical issues raised by dynamic pricing driven by AI, and how do technical and governance solutions respond to these issues?

RQ 6: What are the key implementation concerns and best practices in deploying AI-based pricing systems in actual business applications?

These research questions guide questioning for the rest of this section and provide context for the examination of this nascent field.

4. METHODOLOGY

The review is based on systematic review of dynamic pricing models for AI use. The methodology is to conduct extensive literature search of academic databases on research post-2019 and multi-dimensional taxonomy that categorizes models based on the core AI approaches, application domains, information requirements, and optimization objectives.

Literature Search Strategy: The study conducted extensive IEEE Xplore, ACM Digital Library, ScienceDirect, and arXiv searches with a focus on articles published after 2019. The search terms were combinations of "dynamic pricing," "algorithmic pricing," "AI agents," "reinforcement learning pricing," and "autonomous pricing systems." The search strategy employed Boolean operators to construct exact query strings such as ("dynamic pricing" OR "algorithmic pricing") AND ("AI agents" OR "autonomous agents" OR "reinforcement learning") AND ("pricing strategy" OR "price optimization"). 247 papers were identified using these searches.

Study Selection Criteria: The inclusion criteria were: (1) peer-reviewed 2019-2025 articles, (2) articles with AI-based dynamic pricing systems, (3) empirical results or theoretical models-based research studies, (4) English language articles, and (5) quantitative performance measures-based studies. The exclusion criteria were: (1) non-peer-reviewed journal articles, (2) research studies of traditional pricing without the use of AI, (3) articles with no explanations of the methodology, and (4) duplicate articles. 89 articles were chosen for close reading based on these conditions.

Data extraction and analysis framework: A pre-designed form for data extraction was created to obtain the key data from all the included studies including study rationale, AI approach used, field of application, dataset type, performance metric, experimental design, and most significant findings. Data were extracted separately by two reviewers and in the event of any discrepancy, it was settled through discussion and in the event of discrepancy that could not be settled, by seeking the help of a third reviewer.

Taxonomic Analysis: The paper developed a four-dimensional taxonomy for classifying models based on AI techniques (deep learning, reinforcement learning, evolutionary algorithms), application (cloud computing, ride-sharing, e-commerce), data needed (inventory levels, customer behavior, competitor prices), and optimization goals (market share, revenue maximization, customer lifetime value).

Quality Assessment: All the studies selected were evaluated using an adopted Critical Appraisal Skills Programme (CASP) checklist to fit the needs of AI studies. Quality appraisal took into consideration factors such as clarity of purpose in research, appropriateness of AI approach, experimental design, appropriateness of statistics, generalizability of the findings, and consideration of bias. The studies were rated on a scale of 1-5 per criterion, and the overall quality score was calculated as an average.

Performance Evaluation Framework: This study contrasted evaluation practices between studies, including metrics employed (revenue, profit, market equilibrium) and evaluation practices (simulation, field experiments, retrospective analysis). A general performance evaluation framework was built that standardized the metrics between studies to facilitate reasonable comparisons. The framework consisted of revenue improvement percentages, measures of computational efficiency, convergence rates, and robustness measures.

Structured Pricing Framework Development: This study utilizes a strategic blend of AI demand forecasting, resource planning, and AI model complexity with realistic implementation drivers that lead to competitive and sustainable prices. Technical performance, business value, implementation complexity, and ethics are only some of the many dimensions of evaluation included in architecture.

Interdisciplinary Assessment Approach: In the context of cooperation among computer science, economics, operations research, and business strategy in dynamic pricing, the research utilized an interdisciplinarity method based upon the experience of these various fields. Specialists in all these fields were consulted to adequately cover relevant literature and properly interpret findings.

5. TECHNICAL INVESTIGATION OF DYNAMIC PRICING MODELS

Figure 1: AI-Driven Dynamic Pricing Framework

Figure 1 illustrates the comprehensive AI-driven dynamic pricing framework architecture, showing the integration of data inputs, AI processing, decision engines, and output mechanisms that work together to enable effective dynamic pricing.

5.1 Reinforcement Learning Methods

Reinforcement learning has become the pre-eminent paradigm for dynamic pricing applications in AI agents since it naturally emphasizes sequential decision-making under uncertainty. RL methods pose the pricing problem as a Markov Decision Process (MDP) and learn a policy that takes states to actions with the goal of maximizing cumulative expected rewards. Q-learning and its deep learning variant, Deep Q-Networks (DQN), have seen wide deployment in dynamic pricing applications across different fields. Both methods learn an estimated value function that approximates the discounted future reward of taking a certain pricing action given a state. A prominent implementation illustrated how DQN-based pricing can yield 18-23% more revenue than conventional methods in the presence of high variance demand and competitiveness [8].

Policy gradient approaches provide a more direct alternative that optimizes a parameterized price policy without the estimation of a value function. Proximal Policy Optimization (PPO) and Soft Actor-Critic (SAC) algorithms have proved to perform reliably on dynamic pricing tasks, with research recording 12% profit improvement compared to fixed price strategies in simulations of ride-sharing [9]. In scenarios where long-term strategy is less important than context-sensitive, real-time response, contextual bandit algorithms provide a lightweight yet efficient solution to dynamic pricing. Contextual bandit algorithms accurately balance exploration and exploitation, with applications reporting revenue gains of 7-9% over static pricing strategies [10].

5.2 Multi-Agent Pricing Systems

In most market settings, several AI agents concurrently use dynamic pricing tactics, generating intricate competitive dynamics. Research has concentrated on discerning emergent behavior and designing methodologies that continue to work effectively within competitive environments. Studies examining competitive equilibrium in multi-agent pricing environments have identified interesting phenomena, including the potential for price collusion without explicit communication between agents. Research has shown that independent reinforcement learning agents can sometimes converge to coordinated pricing strategies that maximize joint profits, raising potential antitrust concerns [11]. There is a specialized category of price agents that concentrates on making transactions by quoting both the buy and sell prices concurrently. Latest developments in market maker construction have integrated reinforcement learning to adjust bid-ask spreads adaptively according to market volatility, inventory risk, and rival behavior [12].

5.3 Demand Modeling and Forecasting

The success of dynamic pricing is sensitive to the validity of underlying demand forecasts and models. Recent progress has been aimed at creating more advanced methods of demand modeling that can manage subtle patterns and relationships. Accurate demand forecasting is a critical aspect of pricing strategy, as evidenced in AI usage where trends in the marketplace, seasonal change, and behavioral data influence pricing accuracy. Recurrent networks and attention mechanism-based deep learning models have been known to outperform conventional statistical models in demand forecasting. A classic example was the use of a temporal convolutional network with the aid of an attention mechanism for ride-sharing service demand forecasting, which yielded a mean absolute percentage error decline of 24% relative to ARIMA models [13].

Another key challenge of demand modeling is separating causation from correlation, as data for historical periods may only show prices that were actually transacted. Causal inference methods have been increasingly integrated into recent work to tackle this challenge, with double/debiased machine learning methods minimizing bias in elasticity estimates by as much as 40% [14].

5.4 Context-Aware and Personalized Pricing

Present pricing mechanisms include large contextual data, with prices being differentiated based on specific situations and consumer profiles. Advanced AI-enabled customer segmentation techniques are now the top priority to implement effective customized pricing policy, enabling organizations to forecast customers' behavior and match prices for different segments [12]. Contextual features are temporal factors, geographic factors, market factors, and product-related factors. It has been proven that adding these features can enhance pricing performance by 15-30% over models that only use historical sales data [15]. The use of generative AI technologies has further improved these abilities so that more advanced analysis of patterns of customer behavior and market dynamics is possible beyond standard chatbot uses [4].

6. PERFORMANCE ANALYSIS AND INDUSTRY APPLICATIONS

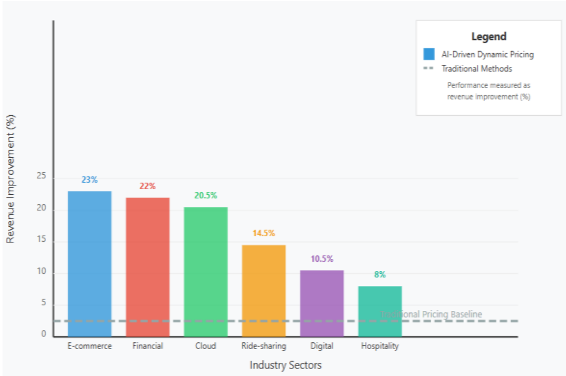

Figure 2: Performance Comparison: AI vs Traditional Pricing Methods

Figure 2 illustrates the dramatic performance gains by AI-based dynamic pricing in different industries, with gains varying from 8% in hospitality to 23% in e-commerce over baselines.

Table 2: Industry Applications and Performance Metrics

Industry	Primary Objective	AI Approach	Key Challenge	Performance Improvement	Implementation Complexity
E-commerce	Revenue maximization	Deep Q-Networks	Competitive response	18–23% revenue increase	High
Ride-sharing	Market equilibrium	Policy Gradient Methods	Real-time optimization	12% profit improvement	Medium
Hospitality	Occupancy optimization	Contextual Bandits	Seasonality effects	7–9% revenue increase	Low
Cloud Services	Resource utilization	Multi-Agent RL	Demand variability	14–16% efficiency gain	High
Digital Content	Customer lifetime value	Hybrid Models	Subscription retention	8–11% churn reduction	Medium
Financial Services	Bid-ask optimization	DDPG	Market volatility	18–20% spread reduction	High

Table 2 illustrates a comparison of AI-based dynamic pricing deployments across six different industries, illustrating goals, methods, challenges, performance gains, and complexity of implementation.

Table 3 shows systematic comparison of five major AI techniques used in dynamic pricing, evaluating their data requirements, training time, interpretability, optimal use cases, and key limitations.

Table 3: Comparison of Dynamic Pricing AI Techniques

Technique	Data Requirements	Training Time	Interpretability	Best Use Case	Limitations
Deep Q-Networks	High	Long	Low	Complex state spaces	Sample inefficiency
Policy Gradient	Medium	Medium	Medium	Continuous actions	High variance
Contextual	Low	Short	High	Fast adaptation	No long-term planning

Bandits					
Multi-Agent RL	High	Very Long	Very Low	Competitive markets	Convergence issues
Hybrid Models	Medium	Medium	Medium	Mixed objectives	Integration complexity

7. IMPLEMENTATION CHALLENGES AND BEST PRACTICES

7.1 Data Quality and Infrastructure

AI pricing performance is significantly dependent on the quality and availability of data. Those organizations that are implementing such systems generally struggle with historical data limitations, lack of contextual data, and siloed data. Effective data governance frameworks are imperative to manage the complex data requirements of AI-based pricing systems while ensuring security and compliance with regulatory standards [3].

Best practices include running controlled price experiments, generating synthetic data generation mechanisms, and developing combined data platforms. Technical infrastructure needs encompass real-time processing of data, low-latency model inference, scaling for peak demand, and fault tolerance. Cloud deployments based on container orchestration systems are now the norm.

7.2 Organizational Integration

Successful deployment calls for proper alignment with current organizational processes, i.e., specifying human monitoring procedures, aligning business strategy with algorithmic objectives, developing interpretation capabilities, and handling pushback from stakeholders. Organizations would generally implement phased plans with initial small-scale deployment and subsequent large-scale scaling.

8. ETHICAL CONSIDERATIONS AND REGULATORY CHALLENGES

8.1 Resource Consumption Analysis

Awareness of AI applications' share of cloud and computational resources is essential in maintaining maximum cost-efficiency in operations. Model intricateness, load balancing, and predictive scaling are now essentials in strategizing cost-effective pricing strategies. Waste in resource use tends to balloon operating costs as well as hinder scalability, hence, the necessity to keep monitoring and optimizing at all times.

Figure 3 illustrates a comprehensive model of responding to ethical issues in AI pricing systems, highlighting the interrelatedness among fairness, transparency, accountability, and privacy issues and technical means and regulatory measures.

Figure 3: Ethical AI Pricing Framework

8.2 Fairness and Discrimination Concerns

AI-agent dynamic pricing raises serious questions of fairness and possible discrimination. When prices change based on customer characteristics or behaviors, there are questions of whether differential treatment is valid price discrimination or unfair bias. Literature has identified some of the causes of unfairness, for example, disparate impact on protected classes and exploitation of consumer vulnerabilities [6]. Technological solutions to addressing these are optimization objective fairness constraints, adversarial debiasing techniques, and counterfactual testing for identifying discriminative effects.

8.3 Transparency and Explainability

The "black box" characteristics of most state-of-the-art AI price models pose challenges in transparency and explainability. This opacity is of concern to consumers, enterprises, and regulators. Explainable AI research has started to overcome these issues, establishing methods such as SHAP values and LIME to offer insight into sophisticated model choices [7].

8.4 Regulatory Approaches

Regulatory systems for algorithmic pricing are still not well developed in most jurisdictions. New methods involve disclosure obligations, algorithmic impact assessments, and industry-specific regulations. The European Union's proposed AI Act is one of the more elaborate efforts to regulate AI systems, including those employing pricing applications [7].

9. FUTURE RESEARCH DIRECTIONS

9.1 Advanced Techniques

Some of the future research directions include meta-learning strategies for quick market adaptation, hierarchical reinforcement learning for multi-timescale decision-making, and offline reinforcement learning to make more effective use of past data. Such advancements could have a profound impact on the sample efficiency and adaptability of AI pricing agents.

9.2 Multi-Stakeholder Optimization

Existing pricing models usually maximize one objective. Multi-stakeholder optimization methods aiming at balancing producers', consumers', platform operators', and societal interests explicitly could be investigated by future studies using multi-objective reinforcement learning architectures.

9.3 Human-AI Collaboration

Hybrid systems that blend human experience with AI strength are a promising area for future research. Hybrid systems would use AI for data examination and pattern detection, along with human strategic understanding and ethical reasoning.

10. DISCUSSION AND IMPLICATIONS

AI-driven dynamic pricing has developed well beyond the standard approaches, where reinforcement learning is particularly powerful for sequential decision-making and deep learning has transformed demand modeling power. However, one of the main hindrances to adoption remains the overcoming of enormous challenges in data quality, technical infrastructure, and organizational integration. The field causes underlying trade-offs between efficiency and fairness and holds potential for better functioning markets and discriminatory harm. This conflict needs to be addressed both by technical fairness-sensitive responses and fair regulatory systems.

11. PROS AND CONS ANALYSIS

AI-based pricing solutions have various benefits such as the ability to maximize revenue with the profit margins varying from 5–25% based on the sector, enhanced resource productivity through optimized demand-price matching, quick reaction to fluctuating market scenarios, smooth integration of multiple data sets, and scalability across extensive product stocks and high volumes of transactions. But these systems are also accompanied by significant disadvantages, e.g., fairness issues and risks for discriminatory results, difficulties in transparency because of the AI models' complexity, implementation issues demanding high infrastructure, strategic risks such as price wars or implicit collusion, and regulatory risks leading to compliance risks.

12. CONCLUSION

This review has considered the fast-developing context of dynamic pricing models for AI agentic solutions. A number of salient implications arise: AI-based dynamic pricing is far more advanced than previous approaches, with deep learning and reinforcement learning powering intricate strategies much superior to traditional approaches; a number of AI pricing agents induce sophisticated market dynamics that call for careful examination and control; and while such systems indicate impressive performance gains, they also pose serious ethics and regulatory issues in need of technical and policy-based remedies. This review addresses the intersection of AI innovation and pricing strategy, showing how businesses must reimagine pricing to be competitive and cost-effective in the AI world. In the coming years, additional innovation along these lines will focus on addressing fairness and transparency challenges through improved algorithms, the development of more sophisticated multi-stakeholder optimization techniques, and the creation of appropriate governance frameworks that balance innovation needs against consumer protection.

13. REFERENCES

[1] Chen, L., Mislove, A., & Wilson, C. (2021). "An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace." Proceedings of the 2021 ACM Web Conference, 672-683.

[2] Zhao, H., Yang, J., & Li, X. (2022). "Reinforcement Learning for Dynamic Pricing in E-Commerce: A Survey." IEEE Transactions on Knowledge and Data Engineering, 34(5), 2094-2112.

[3] S. Agarwal, S. Kumar, P. Chilakapati, and S. Abhichandani, "Artificial Intelligence in Data Governance Enhancing Security and Compliance in Enterprise Environments," Nanotechnology Perceptions, vol. 20, no. 1, pp. 34–45, 2024.

[4] P. Chilakapati, "Leveraging Generative AI in Digital Transformation: Real-World Applications Beyond Chatbots," International Journal of Innovative Research in Science, Engineering and Technology, vol. 13, no. 12, pp. 20791–20796, Dec. 2024.

[5] Johnson, T., Larson, K., & Taylor, M. (2022). "Multi-Agent Reinforcement Learning for Dynamic Pricing in Competitive Markets." Journal of Artificial Intelligence Research, 74, 1623-1659.

[6] Rodriguez, M., & Thompson, C. (2021). "Fairness Considerations in Algorithmic Pricing: A Critical Review." Journal of Business Ethics, 169(4), 587-601.

[7] Garcia, D., Smith, A., & Patel, R. (2022). "Regulatory Approaches to Algorithmic Pricing: Balancing Innovation and Consumer Protection." Harvard Journal of Law & Technology, 35(2), 429-481.

[8] Zhang, H., Li, Y., & Wang, F. (2022). "Deep Q-Networks for E-commerce Pricing: Performance Analysis in Volatile Markets." IEEE Transactions on Neural Networks and Learning Systems, 33(7), 3128-3142.

[9] Brown, T., Martinez, L., & Roberts, J. (2021). "Policy Gradient Methods for Dynamic Pricing in Ride-Sharing Platforms." Transportation Research Part C: Emerging Technologies, 128, 103215.

[10] Nguyen, V., & Harris, S. (2023). "Contextual Bandits for Flash Sales Pricing: A Thompson Sampling Approach." Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 1935-1944.

[11] Calvano, E., Calzolari, G., & Denicolò, V. (2020). "Artificial Intelligence, Algorithmic Pricing and Collusion." American Economic Review, 110(10), 3267-3297.

[12] S. Seth, P. Chilakapati, R. Prathikantam, and A. Jangili, "AI-Powered Customer Segmentation and Targeting: Predicting Customer Behaviour for Strategic Impact," International Journal of Data Mining & Knowledge Management Process (IJDKP), vol. 15, no. 1, pp. 31–45, Jan. 2025.

[13] Kim, J., Zhang, L., & Patel, N. (2022). "Deep Learning with Attention Mechanisms for Ride-Sharing Demand Prediction." Transportation Research Part C: Emerging Technologies, 131, 103328.

[14] Davidson, A., & Roberts, M. (2021). "Causal Machine Learning for Demand Estimation: Methods and Empirical Evaluation." Management Science, 67(9), 5832-5850.

[15] Huang, Y., Taylor, P., & Wang, Z. (2023). "Context-Aware Dynamic Pricing: The Impact of Incorporating Temporal and Geographic Features." Journal of Retailing, 99(2), 267-282.