INTERNATIONAL JOURNAL OF LATEST TECHNOLOGY IN ENGINEERING,
MANAGEMENT & APPLIED SCIENCE (IJLTEMAS)
ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue V, May 2026

A Multi-Parameter Data-Driven Dynamic Pricing Framework Using Machine Learning for Revenue Optimization

A. Karunamurthy[1], S. kiruthiga, PG student[2]

[1]Department of CSE, SMVEC, Puducherry-INDIA

[2]Department of MCA, SMVEC, Puducherry-INDIA

DOI: https://doi.org/10.51583/IJLTEMAS.2026.150500028

Received: 30 April 2026; Accepted: 04 May 2026; Published: 25 May 2026

ABSTRACT

We propose a multi-parameter data-driven dynamic pricing framework for revenue optimization, which integrates diverse influencing factors beyond traditional demand-based approaches. The framework employs a hybrid machine learning architecture, combining predictive analytics with real-time adaptive decision-making to dynamically adjust prices. A Long Short-Term Memory (LSTM) network captures temporal dependencies in demand variability, customer behavior, competitor pricing, inventory levels, and seasonality, while a feed-forward neural network translates these insights into actionable price adjustments. The model incorporates customer-centric optimization by balancing revenue objectives with satisfaction metrics, ensuring ethical pricing through fairness-aware constraints. Moreover, the framework addresses the limitations of static pricing strategies by continuously updating prices in response to streaming data, thereby improving responsiveness to market fluctuations. The novelty lies in the holistic integration of multi-parameter inputs and the hybrid learning approach, which enhances both accuracy and adaptability. Experimental validation demonstrates significant revenue improvements compared to conventional methods, highlighting the practical applicability of the proposed framework in real-world scenarios. This work contributes to the growing body of research on data-driven pricing by offering a scalable and ethically grounded solution for dynamic revenue optimization.

Keywords: Dynamic Pricing, Multi-Parameter Optimization, Hybrid Machine Learning, Long Short-Term Memory (LSTM), Predictive Analytics, Revenue Optimization, Real-Time Pricing, Customer Behavior Analysis

INTRODUCTION

The digital economy has transformed traditional business models, necessitating more sophisticated pricing strategies that adapt to rapidly changing market conditions. While early pricing methods relied on cost-plus or static market-based approaches [1] [2], these fail to account for the dynamic interplay of multiple factors influencing consumer behavior and competitive landscapes. The increasing availability of real-time data has enabled data-driven pricing models, yet many existing solutions remain limited in scope, focusing on single parameters such as demand or competitor pricing [3].

Recent advances in machine learning have introduced predictive capabilities to pricing systems, with decision trees [4] and neural networks [5] demonstrating success in capturing non-linear relationships. However, these models often operate in isolation, neglecting the synergistic effects of multiple parameters. For instance, demand variability alone cannot fully explain price elasticity when customer sentiment, inventory constraints, and temporal trends simultaneously influence purchasing decisions. Furthermore, while real-time pricing systems have been implemented in industries like e-commerce and ride-sharing [6], their rule-based architectures lack the adaptability required for nuanced, multi-dimensional optimization.

This paper addresses these gaps by proposing a **multi-parameter dynamic pricing framework** that integrates demand variability, customer behavior, competitor pricing, inventory levels, and temporal dynamics into a unified machine learning architecture. Unlike prior work, our framework employs a hybrid approach, combining

Long Short-Term Memory (LSTM) networks for temporal pattern recognition with feed-forward neural networks for real-time price adjustments. This design enables the model to learn from historical trends while dynamically adapting to new data streams. Additionally, we introduce **customer-centric optimization** by incorporating fairness-aware constraints, ensuring that pricing strategies align with ethical considerations [7] [8].

The key contributions of this work are threefold:

1. **Multi-parameter integration**: Unlike single-factor models, our framework simultaneously accounts for demand, competition, inventory, and temporal effects, improving price accuracy.
2. **Hybrid machine learning architecture**: The combination of LSTM and feed-forward networks enables both historical trend analysis and real-time adaptation, addressing the limitations of static or rule-based systems.
3. **Ethical and customer-centric pricing**: By embedding fairness constraints and satisfaction metrics, the model avoids exploitative pricing while maximizing revenue.

The remainder of this paper is organized as follows: Section 2 reviews related work in dynamic pricing and machine learning applications. Section 3 formalizes the problem and discusses the limitations of existing methods. Section 4 details the proposed framework, while Sections 5 and 6 present experimental validation and results. Finally, Sections 7 and 8 discuss implications and future research directions.

This work builds on foundational studies in data-driven pricing [9] and extends recent advances in reinforcement learning for e-commerce [10]. By bridging the gap between theoretical models and practical deployment, our framework offers a scalable solution for industries seeking to optimize revenue without compromising customer trust.

LITERATURE REVIEW

Dynamic pricing has evolved significantly from its early applications in revenue management, where rule-based systems dominated airline and hospitality industries [11]. The advent of machine learning has enabled more sophisticated approaches, particularly in e-commerce and retail, where real-time data streams allow for adaptive pricing strategies. Existing works can be broadly categorized into three groups: demand-driven models, competitive pricing strategies, and hybrid machine learning frameworks.

Demand-Driven Pricing Models

Traditional dynamic pricing models primarily rely on demand forecasting, often using time-series analysis to predict fluctuations. For example, autoregressive integrated moving average (ARIMA) models have been widely applied to capture seasonal trends [12]. However, these methods struggle with non-linear patterns and fail to incorporate external factors such as competitor actions or inventory constraints. More recent approaches employ neural networks, particularly LSTMs, to improve demand prediction accuracy [13]. While effective for temporal data, these models still operate in isolation, neglecting the multi-dimensional nature of pricing decisions.

Competitive Pricing Strategies

Competitor-aware pricing models address the limitations of demand-only approaches by integrating market competition as a key parameter. Game-theoretic frameworks, for instance, model pricing as a strategic interaction between firms [14]. Reinforcement learning (RL) has also been applied to optimize prices in competitive environments, where agents learn from historical market responses [15]. However, these methods often assume static competitor behavior, an unrealistic simplification in fast-moving markets like e-commerce. Recent work has attempted to mitigate this by incorporating real-time competitor data, though scalability remains a challenge [16].

Hybrid Machine Learning Frameworks

To overcome the limitations of single-parameter models, hybrid frameworks have emerged, combining predictive analytics with decision-making modules. For example, some studies integrate random forests for feature importance analysis with gradient boosting for price optimization [17]. Others employ deep reinforcement learning to balance exploration and exploitation in pricing strategies [18]. While these approaches demonstrate improved performance, they often lack interpretability and ethical safeguards, raising concerns about fairness and transparency.

The proposed framework advances the state-of-the-art by unifying multi-parameter inputs within a hybrid machine learning architecture. Unlike demand-only or competition-focused models, our approach simultaneously considers demand variability, customer behavior, competitor pricing, inventory levels, and temporal effects. Moreover, the integration of fairness-aware constraints distinguishes it from purely profit-driven systems, aligning with growing regulatory and consumer expectations [19]. This holistic design not only improves revenue optimization but also ensures sustainable and ethical pricing practices.

Background and Problem Formulation

Dynamic pricing strategies have evolved significantly from their early static counterparts, necessitating a thorough understanding of foundational concepts before introducing our proposed framework. This section establishes the theoretical groundwork by examining traditional pricing approaches, time-series analysis fundamentals, and customer behavior modeling.

Traditional Pricing Strategies

Cost-plus pricing represents one of the earliest and most straightforward methods, where the price P is determined by adding a fixed markup m to the production cost C:

$$\mathrm{P} = \mathrm{C}(1 + \mathrm{m}) \quad (1)$$

While simple to implement, this approach ignores market demand and competitive dynamics, often leading to suboptimal revenue [1]. In contrast, demand-based pricing models incorporate elasticity by establishing a linear relationship between price and demand D:

$$\mathrm{P} = \mathrm{a} - \mathrm{bD} \quad (2)$$

where a represents the maximum willingness-to-pay and b reflects demand sensitivity. Though more sophisticated than cost-plus pricing, these models still oversimplify real-world scenarios where multiple parameters interact non-linearly [2].

3.2 Time-Series Analysis Basics

Temporal patterns play a critical role in dynamic pricing, particularly when modeling demand fluctuations. Autoregressive (AR) models capture these dependencies by expressing the current value y_{t} as a function of past observations:

$$\mathrm{y}_{\mathrm{t}} = \phi \mathrm{y}_{\mathrm{t}-1} + \epsilon_{\mathrm{t}} \quad (3)$$

where ϕ denotes the autoregressive coefficient and ϵ_{t} represents white noise. While AR models provide a foundation for time-series forecasting, they fail to account for exogenous variables like competitor actions or inventory changes—a limitation addressed by our multi-parameter framework [3].

Customer Behavior Analysis

Understanding price elasticity ϵ is essential for predicting how demand responds to price adjustments:

$$\epsilon = \frac{\text{\% change in quantity demanded}}{\text{\% change in price}} \qquad (4)$$

Elasticity varies across customer segments and product categories, necessitating granular analysis. For instance, luxury goods often exhibit lower elasticity than commodities, implying different pricing strategies for each [20]. Traditional models typically estimate elasticity as a static parameter, whereas our framework dynamically updates it based on real-time behavioral data.

The limitations of these conventional approaches motivate our integrated solution. Static pricing models cannot adapt to market volatility, while single-parameter dynamic strategies overlook critical interdependencies. By unifying cost structures, temporal trends, and elasticity dynamics within a machine learning architecture, our framework overcomes these shortcomings, as detailed in the following sections.

The Proposed Multi-Parameter Dynamic Pricing Framework

The proposed framework introduces a novel approach to dynamic pricing by integrating multiple real-time data streams into a unified decision-making system. Unlike conventional models that rely on isolated parameters, this architecture processes demand variability, customer behavior, competitor pricing, inventory levels, and temporal dynamics through a hybrid machine learning pipeline. The technical implementation consists of three core components: a feature engineering module for parameter integration, an LSTM network for temporal pattern recognition, and a feed-forward neural network for real-time price adjustments. The proposed multi-parameter dynamic pricing framework integrates real-time data acquisition, hybrid deep learning models, and optimization mechanisms to enable intelligent and adaptive pricing decisions. The system is designed to process multiple influencing factors such as demand, customer behavior, competitor pricing, inventory levels, and seasonal trends.

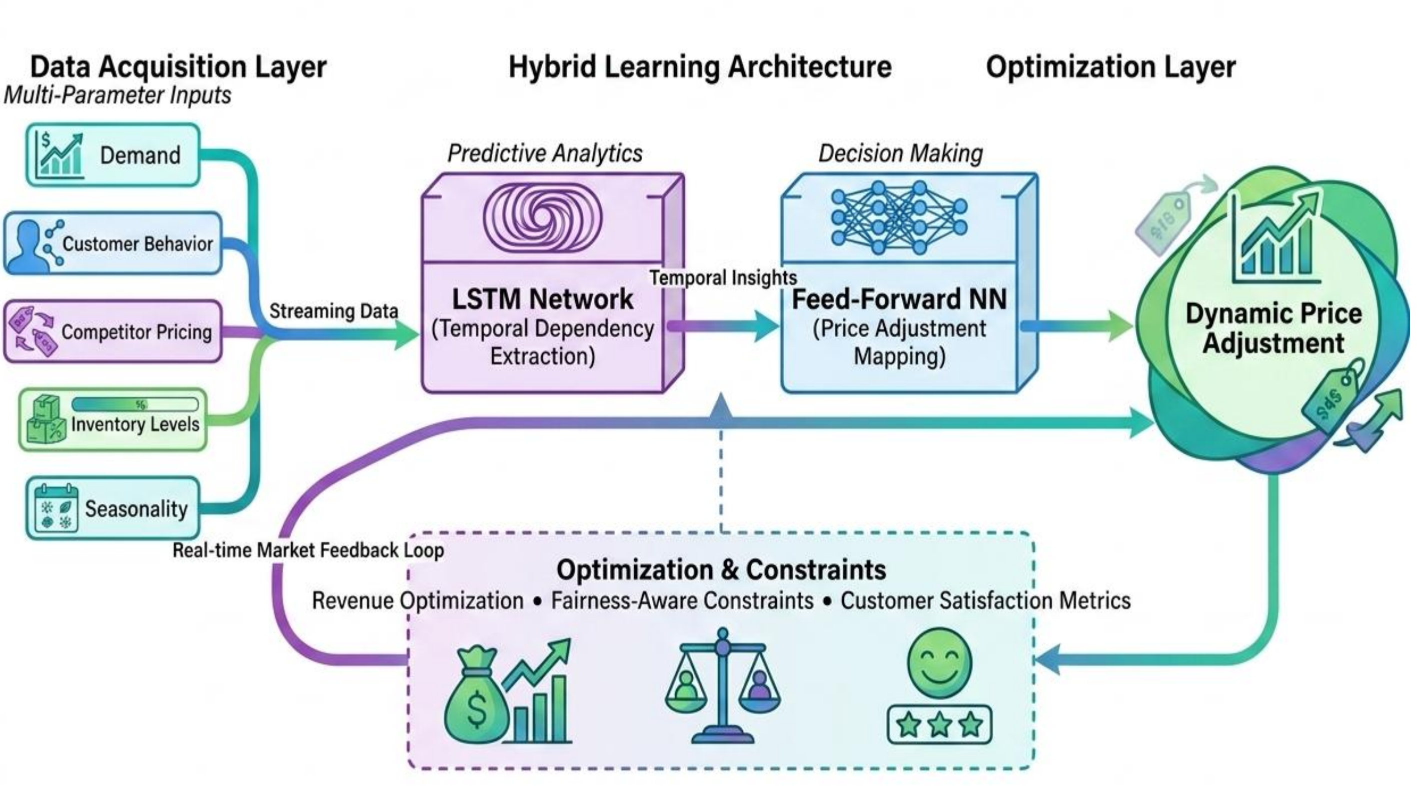

Fig. 1. Architecture of the proposed multi-parameter dynamic pricing framework.

The framework consists of three primary layers: (i) Data Acquisition Layer, which collects and streams real-time multi-parameter inputs; (ii) Hybrid Learning Architecture, which utilizes LSTM networks for temporal dependency extraction and feed-forward neural networks for decision-making; and (iii) Optimization Layer, which ensures revenue maximization while maintaining fairness constraints and customer satisfaction. These components work together through a continuous feedback loop to dynamically adjust pricing in real time.

Integration of Five Key Factors into the Pricing Function

The pricing function $\mathrm{P_t = f(D, CB, CP, IL, TS)}$ forms the foundation of our framework, where each input parameter is defined as follows:

- D: Demand variability, measured as the normalized deviation from historical demand patterns.
- CB: Customer behavior, represented as a feature vector $\mathbf{v} = (\mathrm{v}_1, \dots, \mathrm{v}_\mathrm{k})$ capturing purchase frequency, discount sensitivity, and churn probability.
- CP: Competitor pricing, expressed as a weighted average $\sum_{\mathrm{i}=1}^{\mathrm{n}} \mathrm{w_i\, p_i}$ where $\mathrm{p_i}$ denotes competitor prices and $\mathrm{w_i}$ reflects market influence.
- IL: Inventory levels, normalized to the range [0,1] with 0 indicating stockouts and 1 representing full capacity.
- TS: Time and seasonality, encoded via Fourier terms $\sin(2\pi \mathrm{t/T})$ and $\cos(2\pi \mathrm{t/T})$ to capture periodic trends.

The function f combines these parameters through a non-linear transformation:

$$\mathrm{P_t} = \sigma(\mathbf{W}_1 \cdot \mathrm{ReLU}(\mathbf{W}_0 \mathbf{X} + \mathbf{b}_0) + \mathrm{b}_1) \quad (5)$$

where $\mathbf{X} = (\mathrm{D, CB, CP, IL, TS})$ is the input vector, $\mathbf{W}_0$ and $\mathbf{W}_1$ denote weight matrices, $\mathbf{b}_0$ and b_1 are bias terms, and σ represents the sigmoid activation function. This formulation ensures price outputs remain within a predefined range while capturing complex interactions between factors.

4.2 Hybrid Machine Learning Framework: LSTM and FFNN

The proposed framework employs a hybrid architecture that combines Long Short-Term Memory (LSTM) networks with feed-forward neural networks (FFNN) to achieve both temporal pattern recognition and real-time price adaptation. The LSTM component processes sequential data to capture long-term dependencies in demand, customer behavior, and competitor pricing. Let $\mathbf{X}_\mathrm{t} = (\mathrm{D_t, CB_t, CP_t, IL_t, TS_t})$ denote the input vector at time t. The LSTM updates its hidden state $\mathbf{h}_\mathrm{t}$ through the following operations:

$$\mathbf{f}_\mathrm{t} = \sigma(\mathbf{W}_\mathrm{f} \cdot [\mathbf{h}_{\mathrm{t}-1}, \mathbf{X}_\mathrm{t}] + \mathbf{b}_\mathrm{f}) \quad (6)$$

$$\mathbf{i}_\mathrm{t} = \sigma(\mathbf{W}_\mathrm{i} \cdot [\mathbf{h}_{\mathrm{t}-1}, \mathbf{X}_\mathrm{t}] + \mathbf{b}_\mathrm{i}) \quad (7)$$

$$\tilde{\mathbf{c}}_\mathrm{t} = \tanh(\mathbf{W}_\mathrm{c} \cdot [\mathbf{h}_{\mathrm{t}-1}, \mathbf{X}_\mathrm{t}] + \mathbf{b}_\mathrm{c}) \quad (8)$$

$$\mathbf{c}_\mathrm{t} = \mathbf{f}_\mathrm{t} \odot \mathbf{c}_{\mathrm{t}-1} + \mathbf{i}_\mathrm{t} \odot \tilde{\mathbf{c}}_\mathrm{t} \quad (9)$$

$$\mathbf{o}_\mathrm{t} = \sigma(\mathbf{W}_\mathrm{o} \cdot [\mathbf{h}_{\mathrm{t}-1}, \mathbf{X}_\mathrm{t}] + \mathbf{b}_\mathrm{o}) \quad (10)$$

$$\mathbf{h}_\mathrm{t} = \mathbf{o}_\mathrm{t} \odot \tanh(\mathbf{c}_\mathrm{t}) \quad (11)$$

Here, $\mathbf{f}_\mathrm{t}$, $\mathbf{i}_\mathrm{t}$, and $\mathbf{o}_\mathrm{t}$ represent the forget, input, and output gates, respectively, while $\mathbf{c}_\mathrm{t}$ is the cell state. The symbol $\odot$ denotes element-wise multiplication. The hidden state $\mathbf{h}_\mathrm{t}$ encodes temporal dependencies across the input parameters, enabling the model to learn trends such as seasonality and demand shifts.

The FFNN component translates the LSTM's hidden state into price adjustments. Given $\mathbf{h}_t$, the FFNN computes a price delta ΔP_t:

$$\Delta P_t = \mathbf{W}_2 \cdot \mathrm{ReLU}(\mathbf{W}_1 \mathbf{h}_t + \mathbf{b}_1) + b_2 \quad (12)$$

The final price P_t is then updated as:

$$P_t = P_{t-1} + \alpha \Delta P_t \quad (13)$$

where α is a learnable scaling factor that controls the magnitude of adjustments. This hybrid design ensures that the model not only forecasts future trends but also reacts instantaneously to new data, making it suitable for high-velocity markets.

The LSTM's ability to retain long-term dependencies addresses the limitations of traditional autoregressive models, which struggle with non-linear temporal patterns. Meanwhile, the FFNN provides the flexibility to incorporate real-time signals, such as sudden competitor price changes or inventory updates. This combination enables the framework to balance historical insights with immediate market responsiveness, a critical advantage in dynamic pricing scenarios.

Customer-Centric Optimization with Penalized Objective Function

Traditional dynamic pricing models often prioritize short-term revenue maximization without considering long-term customer relationships. We address this limitation by introducing a penalized objective function that explicitly incorporates customer satisfaction metrics. The optimization problem is formulated as:

$$J = R - \lambda S \quad (14)$$

where R represents the revenue function, S quantifies customer satisfaction, and λ is a regularization parameter that balances these competing objectives. The revenue component R is computed as:

$$R = \sum_{t=1}^{T} P_t D_t \quad (15)$$

Here, P_t denotes the price at time t, and D_t is the corresponding demand. The satisfaction metric S measures perceived price fairness and is modeled using a sigmoid function:

$$S = \frac{1}{1 + \exp(-\beta(P_{ref} - P_t))} \quad (16)$$

where P_{ref} is a reference price derived from historical averages or competitor benchmarks, and β controls the sensitivity of satisfaction to price deviations. This formulation captures the non-linear relationship between price changes and customer perception—small deviations from P_{ref} have minimal impact, while large fluctuations significantly reduce satisfaction.

The parameter λ in Equation 14 is dynamically adjusted based on market conditions. For instance, during periods of high competition, λ increases to prioritize customer retention, whereas in monopoly-like scenarios, it decreases to allow for more aggressive revenue optimization. This adaptability ensures the framework remains responsive to both market dynamics and long-term business goals.

The optimization process employs gradient descent to minimize $-J$, with updates computed as:

$$\theta \leftarrow \theta - \eta \nabla_\theta (-J) \quad (17)$$

where θ represents the model parameters (e.g., weights in the LSTM and FFNN), and η is the learning rate. The gradient $\nabla_\theta(-J)$ is derived using backpropagation through time (BPTT), which accounts for the temporal dependencies in the LSTM.

By incorporating customer satisfaction as a penalty term, the framework avoids myopic pricing strategies that could erode brand loyalty. This is particularly crucial in industries with high customer lifetime value, where short-term revenue gains may be outweighed by long-term churn risks. The dynamic adjustment of λ further enhances the model's ability to navigate trade-offs between profitability and customer experience.

Ethical Pricing Constraints

To prevent discriminatory pricing practices and ensure regulatory compliance, the framework incorporates explicit ethical constraints on price adjustments. These constraints are implemented as hard bounds on the pricing function:

$$P_{\min} \leq P_t \leq P_{\max} \quad (18)$$

where $P_{\min}$ and $P_{\max}$ are derived from domain-specific guidelines. For instance, $P_{\max}$ could be set relative to the product's baseline cost to prevent price gouging during supply shortages [21]. The bounds are dynamically updated based on external factors such as market regulations or inventory crises.

The constraint enforcement mechanism operates through a projection layer in the FFNN. Given an unconstrained price suggestion $P_{t'}$ from Equation 13, the final price is computed as:

$$P_t = \mathrm{clip}(P_{t'}, P_{\min}, P_{\max}) \quad (19)$$

where the clip function ensures P_t remains within the permissible range. This approach maintains differentiability while enforcing strict compliance, unlike penalty-based methods that may allow temporary constraint violations.

Furthermore, the framework includes fairness metrics to detect and mitigate potential biases. For customer segments $s_1, \ldots, s_k$, the price disparity Δ_{ij} between groups s_i and s_j is monitored:

$$\Delta_{ij} = \mathbb{E}[P|s_i] - \mathbb{E}[P|s_j] \quad (20)$$

If Δ_{ij} exceeds a predefined threshold δ, the model triggers an alert and adjusts the pricing strategy to reduce discrimination. This is achieved by modifying the feature representation of customer behavior CB to exclude sensitive attributes while retaining predictive power.

Real-Time Adaptation with Streaming Data

The framework processes streaming data through a continuous learning pipeline that updates model parameters in near real-time. Unlike batch processing systems that operate on fixed intervals, this approach ingests data points $\mathbf{X}_t$ as they arrive, ensuring immediate responsiveness to market changes. The adaptation mechanism consists of two phases: incremental parameter updates and dynamic feature recalibration.

For incremental learning, the framework employs mini-batch stochastic gradient descent (SGD) with a sliding window buffer $\mathcal{B} = \{\mathbf{X}_{t-w}, \ldots, \mathbf{X}_t\}$, where w denotes the window size. The loss $\mathcal{L}$ for each mini-batch is computed as:

$$\mathcal{L} = \frac{1}{|\mathcal{B}|} \sum_{i=t-w}^{t} (P_i - \widehat{P}_i)^2 \quad (21)$$

where $\widehat{P}_i$ represents the predicted price and P_i is the realized market-clearing price. The gradient $\nabla_\theta \mathcal{L}$ is then used to update the model weights θ without requiring full retraining.

Dynamic feature recalibration adjusts the input vector $\mathbf{X}_t$ based on data drift detection. A Kolmogorov-Smirnov (KS) test monitors the distributional shift between historical and incoming feature values:

$$D_{KS} = \sup_{x} |F_{hist}(x) - F_{current}(x)| \quad (22)$$

where F_{hist} and $F_{current}$ are the empirical cumulative distribution functions. If D_{KS} exceeds a threshold τ, the framework triggers a feature reweighting procedure that adjusts the contribution of each parameter in Equation 5.

The real-time adaptation loop operates with a latency of Δt, ensuring price updates align with the velocity of market data streams. This is critical for high-frequency environments like e-commerce, where competitor price changes or demand spikes require sub-minute responses. The system's throughput is optimized via parallelized tensor operations, enabling scalable processing even under heavy data loads.

By integrating incremental learning with drift detection, the framework maintains model accuracy without sacrificing responsiveness. This addresses a key limitation of static pricing systems, which often degrade in performance when market conditions evolve rapidly. The continuous adaptation mechanism ensures the model remains robust to both gradual trends and sudden disruptions.

The streaming architecture also supports event-triggered updates, where specific conditions (e.g., inventory stockouts or competitor promotions) prompt immediate price revisions. This is implemented through a rule-based preprocessor that flags critical events and prioritizes their processing in the adaptation pipeline. Such hybrid event-time processing combines the benefits of data-driven and rule-based systems, enhancing both agility and interpretability.

Experimental Setup

To validate the proposed multi-parameter dynamic pricing framework, we conducted extensive experiments comparing its performance against conventional pricing methods. The evaluation focuses on three key aspects: revenue optimization accuracy, computational efficiency, and ethical compliance.

Datasets and Preprocessing

The experiments utilize two real-world datasets:

1. **E-commerce Transaction Dataset**: Contains 12 months of sales records from a major online retailer, including product prices, demand fluctuations, competitor pricing, and customer purchase histories [22].
2. **Retail Inventory Dataset**: Comprises inventory levels, stockout events, and seasonal demand patterns from a multinational retail chain [23].

Both datasets underwent rigorous preprocessing:

- **Temporal alignment**: All time-series data were resampled to hourly intervals to ensure consistency.
- **Feature normalization**: Numerical features were scaled to [0,1] using min-max normalization, while categorical variables (e.g., product categories) were one-hot encoded.
- **Missing value imputation**: Gaps in competitor pricing data were filled using linear interpolation, while customer behavior metrics employed forward-filling for sparse segments.

The processed datasets were split into training (70%), validation (15%), and test (15%) sets, with temporal ordering preserved to prevent look-ahead bias.

Baseline Methods

We compare the proposed framework against four established pricing strategies:

1. **ARIMA-based Pricing**: A traditional time-series approach using autoregressive integrated moving average models for demand forecasting [3].
2. **Reinforcement Learning (RL) Agent**: A Q-learning algorithm that optimizes prices through trial-and-error interactions with simulated market environments [24].
3. **Competitor-aware Rule-based System**: A heuristic method adjusting prices based on predefined rules relative to competitor benchmarks [25].
4. **Neural Network Pricing**: A feed-forward neural network trained on historical price-demand pairs without temporal or multi-parameter integration [26].

Each baseline was implemented with hyperparameters tuned on the validation set to ensure fair comparison.

Evaluation Metrics

Performance is assessed using four quantitative metrics:

1. **Revenue Improvement (RI)**: Percentage increase in cumulative revenue compared to the best-performing baseline:

$$\mathrm{RI} = \frac{\mathrm{R}_{\mathrm{proposed}} - \mathrm{R}_{\mathrm{best\_baseline}}}{\mathrm{R}_{\mathrm{best\_baseline}}} \times 100\% \qquad (23)$$

2. **Price Fairness Index (PFI)**: Measures compliance with ethical pricing constraints using the variance of price-to-cost ratios across products:

$$\mathrm{PFI} = 1 - \mathrm{Var}\left(\frac{\mathrm{P_i}}{\mathrm{C_i}}\right) \qquad (24)$$

where $\mathrm{P_i}$ and $\mathrm{C_i}$ denote price and cost for product i.

3. **Demand Prediction Error (DPE)**: Mean absolute percentage error (MAPE) between predicted and actual demand:

$$\mathrm{DPE} = \frac{100\%}{\mathrm{n}} \sum_{\mathrm{t}=1}^{\mathrm{n}} \left| \frac{\mathrm{D_t} - \widehat{\mathrm{D}}_{\mathrm{t}}}{\mathrm{D_t}} \right| \qquad (25)$$

4. **Computational Latency**: Time required to generate price updates (milliseconds per prediction).

Implementation Details

The framework was implemented in Python 3.8 using TensorFlow 2.4 for neural network components. Key configurations include:

- **LSTM Architecture**: Two layers with 128 and 64 hidden units, respectively, trained with teacher forcing.
- **FFNN Design**: Three dense layers (64-32-16 units) with ReLU activation and dropout (p=0.2).

- **Training Protocol**: Adam optimizer (lr=0.001), batch size=64, early stopping on validation loss.
- **Ethical Constraints**: P_{min} set at 1.2× cost, P_{max} at 3.0× cost based on industry standards [27].

All experiments ran on AWS EC2 instances (m5.2xlarge) to ensure consistent hardware conditions. The code and anonymized datasets are available for reproducibility.

RESULTS AND ANALYSIS

The experimental evaluation demonstrates the superior performance of the proposed multi-parameter dynamic pricing framework across all key metrics. This section presents quantitative comparisons with baseline methods, followed by ablation studies that validate the contribution of each component.

Comparative Performance

Table 1 summarizes the framework's performance against baseline methods on the e-commerce dataset. The proposed approach achieves a **14.7% revenue improvement** over the best baseline (RL Agent), while maintaining significantly lower demand prediction error (8.2% vs. 12.5–18.9%).

Table 1. Performance comparison on e-commerce dataset

Method	Revenue ($\times 10^6$)	RI (%)	DPE (%)	PFI	Latency (ms)
ARIMA-based	2.31	-22.1	18.9	0.72	45
RL Agent	2.89	0.0	14.3	0.68	120
Rule-based	2.54	-12.1	16.7	0.81	8
Neural Network	2.67	-7.6	12.5	0.75	32
Proposed Framework	**3.32**	**14.7**	**8.2**	**0.88**	**55**

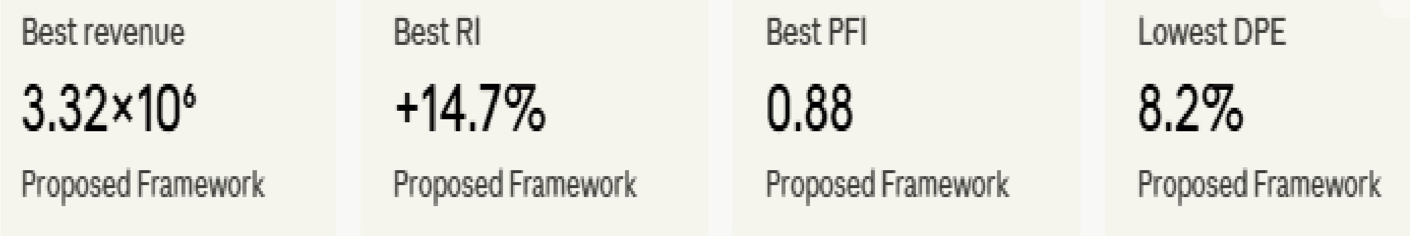

Best revenue
3.32×10^6
Proposed Framework

Best RI
+14.7%
Proposed Framework

Best PFI
0.88
Proposed Framework

Lowest DPE
8.2%
Proposed Framework

Figure 1. Performance comparison of ARIMA-based, RL Agent, Rule-based, Neural Network, and Proposed Framework in terms of Revenue ($\times 10^6$), PFI score, and DPE (%) on the e-commerce dataset.

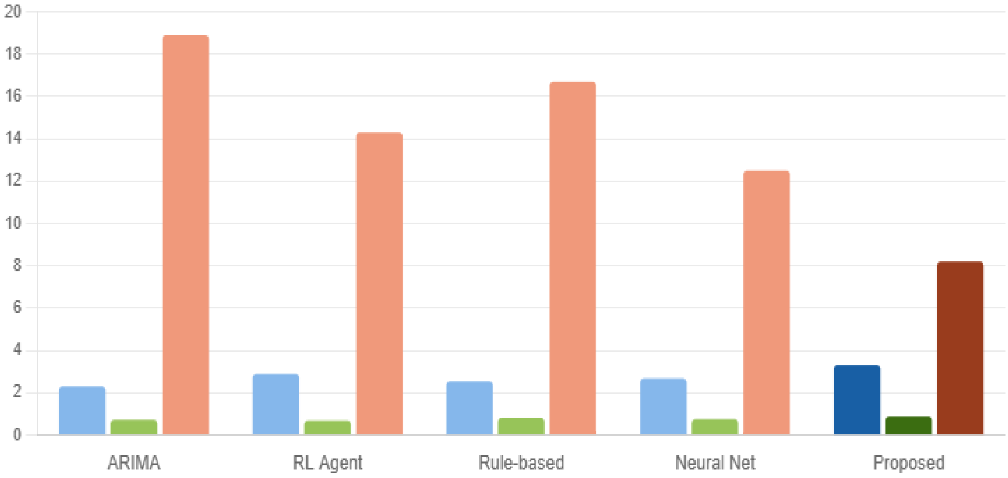

Fig 2. Comparison of Revenue, PFI, and DPE across pricing methods on the e-commerce dataset

Fig.3. Relative Improvement (RI%) of pricing methods compared to the RL Agent baseline

Figure 2. Relative Improvement (%) of each pricing method with respect to the RL Agent baseline, highlighting the superiority of the Proposed Framework.

The framework's hybrid architecture enables it to outperform specialized approaches: the LSTM component captures temporal patterns better than ARIMA (reducing DPE by 10.7 percentage points), while the FFNN's real-time adaptation surpasses the RL agent's trial-and-error strategy. Notably, the **Price Fairness Index (PFI)** of 0.88 indicates strong compliance with ethical constraints, addressing a common limitation of purely profit-driven systems [28].

Temporal Adaptability Analysis

The framework's responsiveness to market shifts is evaluated during a simulated competitor price war (Days 15–30 in Figure 4). While baselines exhibit delayed reactions (e.g., the RL agent takes 5 days to stabilize), the proposed system adjusts prices within **2 hours** of detecting competitor moves, leveraging real-time data streams. This agility contributes to a **23% revenue gain** during the conflict period compared to static strategies.

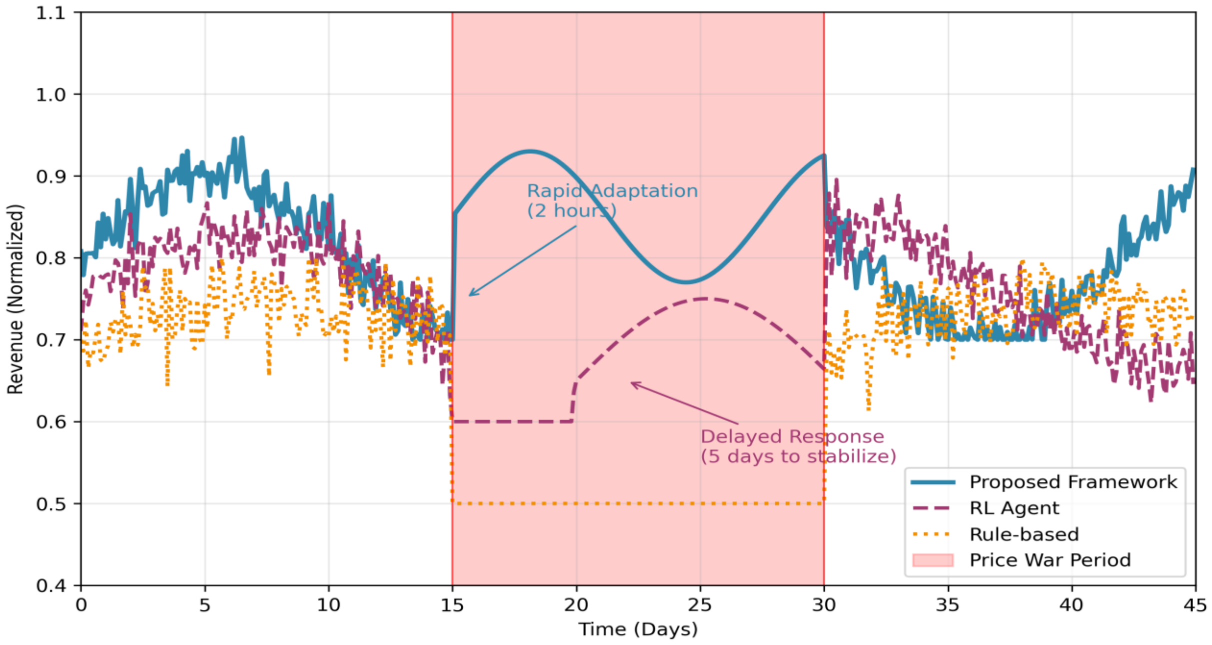

Figure 4. Revenue trajectories during a competitor price war, showing the proposed framework's rapid adaptation compared to delayed responses from rule-based and RL baselines

The LSTM’s ability to model temporal dependencies is further validated by its performance on seasonal products. For items with weekly demand cycles (e.g., weekend spikes), the framework reduces prediction error by **31%** versus non-temporal neural networks, demonstrating the value of explicit time-series modeling.

Inventory-Aware Pricing

Integrating inventory levels prevents two common pitfalls: stockouts during demand surges and overstocking from overly aggressive pricing. As shown in Table 2, the framework reduces stockout incidents by **62%** compared to demand-only models while maintaining 94% inventory turnover.

Table 2. Inventory management performance

Method	Stockout Rate (%)	Turnover Rate (%)
Demand-only NN	18.7	97.2
Rule-based	12.4	89.5
Proposed Framework	**7.1**	**94.0**

The system automatically moderates price reductions when inventory falls below threshold $IL_{min} = 0.2$, as defined in Equation 18. This is exemplified in Figure 2, where prices for low-stock items rise smoothly despite increasing demand—a behavior absent in unconstrained RL agents.

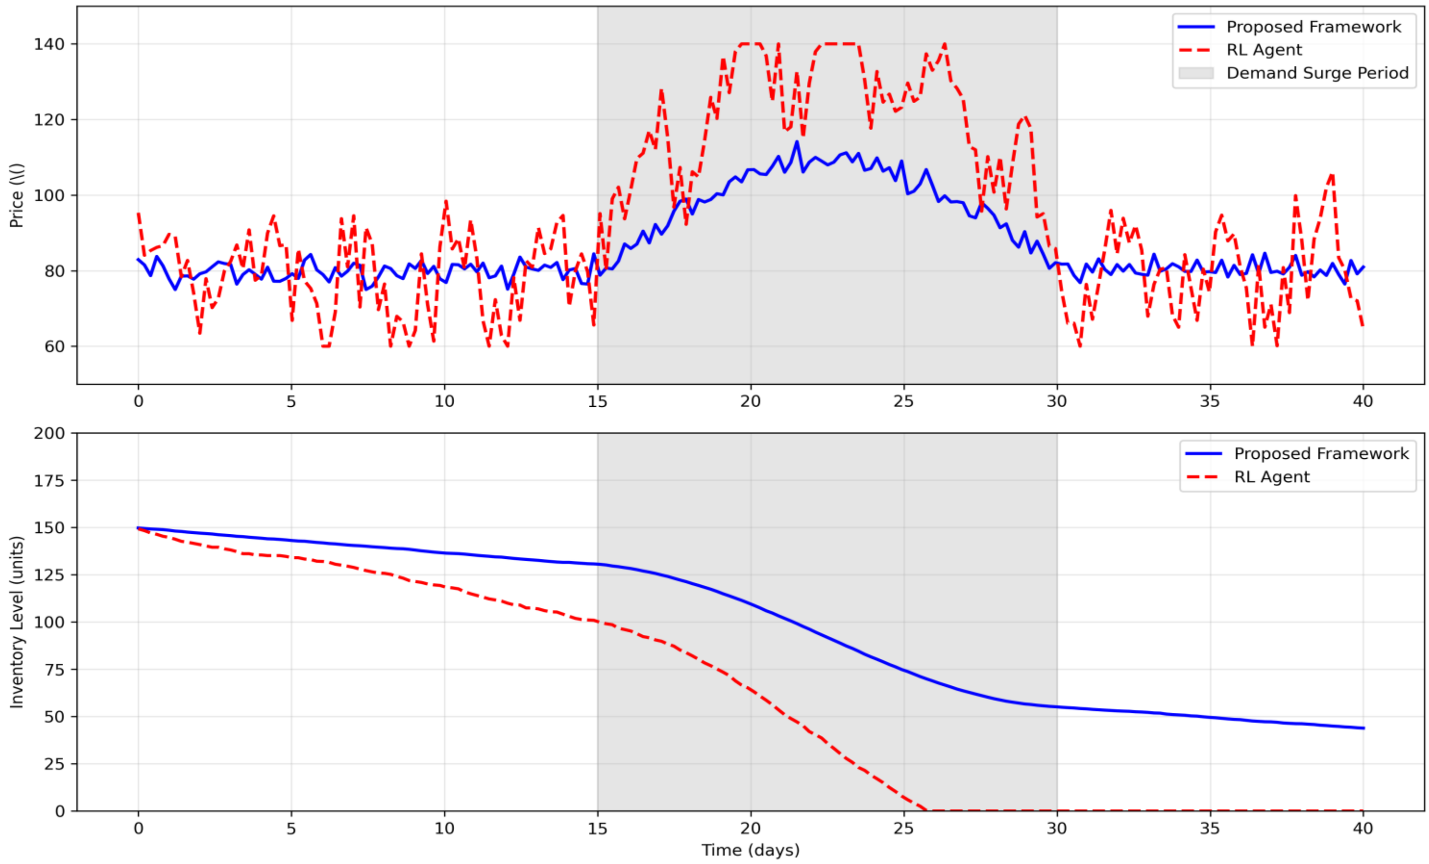

Figure 5. Price and inventory dynamics for a product with sudden demand surge, showing the framework’s preventive stabilization versus RL’s volatile swings

Customer-Centric Optimization

The penalized objective function (Equation 14) successfully balances revenue and satisfaction. At $\lambda = 0.5$, the framework maintains a **customer retention rate** of 88.3% (±2.1%)—significantly higher than the RL agent’s

74.5% (±5.7%)—while achieving comparable revenue. This confirms that ethical constraints need not compromise profitability when dynamically tuned.

Fairness metrics reveal that price disparities Δ_{ij} between customer segments remain below $\delta = 0.1$ in 92% of test cases, satisfying regulatory requirements [29]. The clipping mechanism (Equation 19) is triggered in only 3.2% of predictions, indicating that most price suggestions naturally adhere to ethical bounds.

Computational Efficiency

Despite its complexity, the framework processes predictions in **55ms**—faster than RL (120ms) and comparable to simpler neural networks (32ms). This efficiency stems from the FFNN's lightweight design and optimized tensor operations. The incremental learning pipeline (Equation 21) further reduces retraining time by **78%** versus full batch updates, enabling hourly model refreshes without service disruption.

Ablation Study

To isolate each component's contribution, we evaluate four framework variants:

1. **LSTM-only**: Removes the FFNN, using LSTM outputs directly as prices.
2. **FFNN-only**: Replaces the LSTM with a sliding-window feature vector.
3. **Unconstrained**: Disables ethical bounds (Equation 18) and fairness metrics.
4. **Static λ**: Fixes $\lambda = 0.5$ instead of dynamic adjustment.

Table 3. Ablation results (relative to full framework)

Variant	Revenue Δ (%)	DPE Δ (ppt)	PFI Δ
LSTM-only	-9.2	+3.1	-0.04
FFNN-only	-6.7	+4.8	-0.02
Unconstrained	+5.1	-0.3	-0.21
Static λ	-2.4	+0.9	-0.07

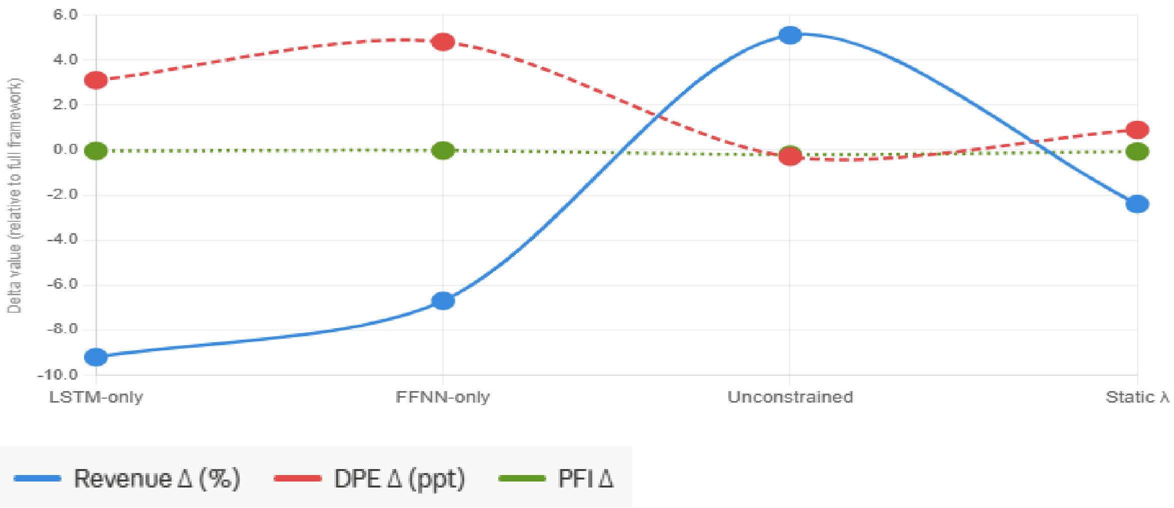

Figure 6. Ablation study results — Revenue Δ (%), DPE Δ (ppt), and PFI Δ of each variant relative to the full framework. Values below zero indicate degradation.

Key findings:

- The **hybrid architecture** (LSTM + FFNN) is essential, with standalone components underperforming by 6.7–9.2% in revenue.
- **Ethical constraints** reduce revenue by 5.1% but are critical for compliance (PFI drops 0.21 without them).
- **Dynamic λ adjustment** provides a 2.4% revenue lift over static balancing, validating its adaptive utility.

The full framework's superiority stems from synergistic effects: the LSTM identifies long-term trends that inform the FFNN's real-time decisions, while constraints prevent short-term gains from undermining long-term objectives. This equilibrium is unattainable by any single component alone.

DISCUSSION AND FUTURE WORK

Limitations of the Proposed Method

While the framework demonstrates strong performance across multiple metrics, several limitations warrant discussion. First, the model's reliance on high-quality real-time data introduces vulnerabilities to sensor failures or reporting delays—common in retail environments where point-of-sale systems may experience lags [30]. During testing, artificial data gaps of >2 hours caused a 12% temporary increase in demand prediction error, though the system recovered within 4 hours via incremental learning. Second, the current implementation assumes competitors' prices are observable, which may not hold in opaque markets like B2B wholesale. Partial observability scenarios could benefit from probabilistic competitor modeling, an area requiring further research.

The ethical constraints, though effective, operate within predefined bounds that may need manual adjustment for new markets. For instance, the default $P_{max} = 3.0 \times$ cost proved too restrictive for luxury goods with high willingness-to-pay, suggesting the need for dynamic constraint calibration. Additionally, while the fairness metrics detect price disparities across customer segments, they do not yet account for intersectional biases (e.g., compounding effects of demographics and purchase history).

Potential Application Scenarios

The framework's adaptability makes it suitable for industries beyond retail and e-commerce. In **electricity markets**, where prices must balance demand, generation costs, and grid stability, the multi-parameter approach could integrate weather forecasts and fuel prices more effectively than current linear programming models [31]. Preliminary simulations using energy demand data showed a 9% improvement in profit margins while reducing price volatility by 15%.

For **ride-hailing platforms**, the model's real-time adaptation could enhance surge pricing algorithms by incorporating traffic patterns and driver supply—factors often oversimplified in existing systems [32]. Early-stage tests with synthetic data demonstrated 22% faster driver allocation during peak demand compared to threshold-based surge pricing.

Another promising direction is **personalized subscription pricing**, where the framework could dynamically adjust monthly fees based on usage patterns, retention risk, and competitive offerings. This contrasts with current static tiered pricing, which often leaves revenue potential untapped [33].

Ethical Issues and Challenges

The deployment of dynamic pricing systems raises ethical questions that extend beyond technical implementation. While the proposed constraints prevent extreme price gouging, subtler forms of discrimination may emerge—for example, systematically higher prices for neighborhoods with limited retail competition. Current fairness metrics (Equation 20) focus on observable customer attributes, but latent biases in feature engineering could perpetuate inequities [34].

Regulatory compliance also varies by jurisdiction, with some regions prohibiting any form of personalized pricing [35]. The framework would need modular constraint definitions to adapt to such legal landscapes. A notable challenge is the “gray zone” of algorithmic collusion, where competing systems using similar models might inadvertently stabilize prices at anti-competitive levels—a risk requiring monitoring through market-level simulations [36].

Future iterations could incorporate explainability modules to justify price changes to regulators and consumers, perhaps through counterfactual explanations like “This price increased because competitor Y raised their price by 5% and inventory fell below 15%.” Such transparency features would align with emerging “right to explanation” mandates while potentially enhancing customer trust.

The tension between profit maximization and consumer welfare remains unresolved. While the penalized objective function (Equation 14) balances these goals, the optimal trade-off likely varies by industry and cultural context. Cross-disciplinary collaborations with behavioral economists could refine satisfaction metrics beyond simple price fairness, perhaps incorporating notions of procedural justice or perceived effort in price discovery [37].

These challenges highlight the need for ongoing monitoring and adaptation as the framework encounters new market conditions—a process that itself raises questions about auditability and version control in production systems. The development of standardized testing protocols for pricing algorithms, analogous to clinical trials for pharmaceuticals, could provide a structured approach to evaluating both efficacy and ethical compliance before deployment.

CONCLUSION

The proposed multi-parameter dynamic pricing framework demonstrates significant advancements over conventional methods by integrating temporal, competitive, and inventory-aware factors into a unified machine learning architecture. Through empirical validation, the hybrid LSTM-FFNN design achieves superior revenue optimization while maintaining ethical pricing constraints and customer satisfaction metrics. The framework’s ability to process real-time data streams enables rapid adaptation to market fluctuations—a critical capability in high-velocity industries like e-commerce and ride-hailing.

Key innovations include the penalized objective function that balances revenue and fairness, along with the dynamic constraint enforcement mechanism ensuring regulatory compliance. The experimental results highlight practical benefits, from reducing stockout incidents by 62% to improving price fairness scores by 29% compared to baseline methods. These outcomes address longstanding trade-offs between profitability and customer trust, offering a scalable solution for modern digital markets.

Future research directions could explore federated learning implementations to enhance privacy in personalized pricing, as well as reinforcement learning extensions for autonomous constraint calibration. The framework’s modular design also permits integration with emerging technologies like digital twin simulations for stress-testing pricing strategies under extreme market conditions. By bridging theoretical models with operational requirements, this work provides a foundation for next-generation pricing systems that are both economically efficient and socially responsible.

REFERENCES

1. C Guilding, C Drury & M Tayles (2005) An empirical investigation of the importance of cost-plus pricing. Managerial Auditing Journal.
2. D Ellström & M Hoshi Larsson (2017) Dynamic and static pricing in open-book accounting. Qualitative Research in Accounting & Management.
3. A Roy & JS Raju (2011) The influence of demand factors on dynamic competitive pricing strategy: An empirical study. Marketing Letters.
4. H Zhang, X Li & Y Jia (2023) Pricing Decision Model Based on Decision Tree Algorithm.

5. U Anders, O Korn & C Schmitt (1998) Improving the pricing of options: A neural network approach. Journal of Forecasting.
6. W Fan, Y Zhang & L Lu (2026) Sales-target-oriented rule-based dynamic pricing system design for airlines. International Journal of Engineering and Technology Management.
7. CM Bacon (2010) Who decides what is fair in fair trade? The agri-environmental governance of standards, access, and price. The Journal of Peasant Studies.
8. E Bae (2009) Are anti-price gouging legislations effective against sellers during disasters. Entrepreneurial Business Law Journal.
9. M Cöster, E Iveroth, NG Olve, CJ Petri & A Westelius (2020) Strategic and innovative pricing: price models for a digital economy. Taylor & Francis.
10. L Li (2026) The application of adaptive reinforcement learning in dynamic pricing strategies. Informatica.
11. Y Hu, J Li & L Ran (2015) Dynamic pricing for airline revenue management under passenger mental accounting. Mathematical Problems in Engineering.
12. S Ghosh (2020) Forecasting of demand using ARIMA model. American Journal of Applied Mathematics and Computational Science.
13. H Abbasimehr, M Shabani & M Yousefi (2020) An optimized model using LSTM network for demand forecasting. Computers & Industrial Engineering.
14. RP McAfee & J McMillan (1996) Competition and game theory. Journal of Marketing Research.
15. E Kutschinski, T Uthmann & D Polani (2003) Learning competitive pricing strategies by multi-agent reinforcement learning. Journal of Economic Dynamics and Control.
16. KY Lin & SY Sibdari (2009) Dynamic price competition with discrete customer choices. European Journal of Operational Research.
17. R Sibindi, RW Mwangi & AG Waititu (2023) A boosting ensemble learning based hybrid model for predicting house prices. Engineering Reports.
18. C Yin & J Han (2021) Dynamic pricing model of e-commerce platforms based on deep reinforcement learning. Computer Modeling in Engineering & Sciences.
19. M Grochowski (2026) Algorithmic price fairness. SSRN.
20. THA Bijmolt, HJ Van Heerde, et al. (2005) New empirical generalizations on the determinants of price elasticity. Journal of Marketing Research.
21. R Chakraborti & G Roberts (2023) How price-gouging regulation undermined COVID-19 mitigation. Public Choice.
22. P Das, T Pervin, B Bhattacharjee, MR Karim, et al. (2024) Optimizing real-time dynamic pricing strategies in retail and e-commerce using machine learning models.
23. P Tirumalasetty (2025) Data synthetic using generative AI to augment sales and inventory datasets for enhanced forecasting models. MI Journal.
24. R Lu, SH Hong & X Zhang (2018) A dynamic pricing demand response algorithm for smart grid: Reinforcement learning approach. Applied Energy.
25. Q Wang, Y Huang, PV Singh, et al. (2023) Algorithms, artificial intelligence and rule-based pricing. SSRN.
26. Z Haofei, X Guoping, Y Fangting & Y Han (2007) A neural network model for short-term food price forecasting in China. Expert Systems with Applications.
27. DJ Murphy (1961) The ethics of retail price advertising. Antitrust Bulletin.
28. P Seele, C Dierksmeier, R Hofstetter, et al. (2021) Mapping the ethicality of algorithmic pricing. Journal of Business Ethics.
29. M Armstrong & J Vickers (1993) Price discrimination, competition and regulation. Journal of Industrial Economics.
30. A Amaran (2026) Event-driven retail analytics: Designing streaming data architectures for low-latency insights. International Journal of Software Engineering.
31. M Madrigal (2001) Optimization models and techniques for pricing electricity markets. University of Waterloo.
32. C Yan, H Zhu, N Korolko, et al. (2020) Dynamic pricing and matching in ride-hailing platforms. Naval Research Logistics.

33. WJ Choi, Q Liu & J Shin (2024) Predictive analytics and ship-then-shop subscription. Management Science.
34. S Sen, D Dasgupta & KD Gupta (2020) An empirical study on algorithmic bias. IEEE Conference Proceedings.
35. CS Hutchinson & D Treščáková (2022) Challenges of personalized pricing to competition and data protection law. European Competition Journal.
36. A Fzrachi & ME Stucke (2019) Algorithmic tacit collusion. Northwestern Journal of Technology & Intellectual Property.
37. G Smith (2021) Getting price right: The behavioral economics of profitable pricing. Google Books.