



Fairlytics: AI-Based Intelligent Dynamic Pricing With Discount Authenticity Validation

Palak Shukla¹, Soukhya Raghavendra Yadawad², Spurthi D S³, Rashmi⁴, Lavanya N L⁵

Student, Dept. of Computer Science and Engineering, East West College of Engineering, Bengaluru, India¹

Student, Dept. of Computer Science and Engineering, East West College of Engineering, Bengaluru, India²

Student, Dept. of Computer Science and Engineering, East West College of Engineering, Bengaluru, India³

Project Guide, Dept. of Computer Science and Engineering, East West College of Engineering, Bengaluru, India⁴

Project Guide, Dept. of Computer Science and Engineering, East West College of Engineering, Bengaluru, India⁵

Abstract: E-commerce platforms increasingly rely on algorithmic pricing to remain competitive, yet two problems persist side by side: prices that fail to track real-time demand, and discount offers that mislead buyers through inflated “strike-through” reference prices, fake festival markdowns, and recycled coupon claims. This paper presents Fairlytics, an AI-based intelligent system that unifies dynamic pricing with automated discount authenticity validation in a single pipeline. Fairlytics estimates a fair, demand-responsive price for a product using historical sales, competitor prices, inventory levels, and seasonal signals, and simultaneously verifies whether an advertised discount is genuine by reconstructing the product's true historical price trajectory and comparing it against the claimed original price. A hybrid ensemble of regression and gradient-boosted tree models drives the pricing engine, while a time-series anomaly detector and rule-based authenticity scorer flag manipulated discounts. The proposed architecture is modular, combining data ingestion, feature engineering, dual prediction engines, and a fairness-scoring dashboard for administrators and consumers. This work is intended to improve pricing transparency, protect consumers from deceptive discounting, and give platform operators a defensible, explainable basis for dynamic pricing decisions.

Keywords: Dynamic Pricing, Discount Authenticity, Fake Discount Detection, Machine Learning, Price Fairness, E-Commerce, Anomaly Detection, Ensemble Learning, Consumer Trust

Highlights

- *AI-based framework for intelligent dynamic pricing and discount validation.*
- *Provides a scalable solution with dashboards for administrators, sellers, and consumers.*
- *Predicts fair product prices using machine learning and historical market data.*
- *Detects fake and misleading discounts through historical price analysis and anomaly detection.*
- *Generates an explainable Fairness Score to improve pricing transparency and consumer trust.*

I. INTRODUCTION

Online retail pricing has fundamentally changed from fixed catalog pricing toward algorithmic and demand-sensitive pricing. E-commerce platforms can adjust product prices several times a day in response to customer demand, competitor prices, inventory availability, seasonal trends, promotional campaigns, product popularity, and broader market conditions. As a result, data-driven pricing has become an important part of modern online retail. Machine learning techniques can analyse historical transactions and market signals to identify patterns and support automated pricing decisions.

Dynamic pricing refers to the continuous or periodic adjustment of a product's price according to changing market conditions. A well-designed dynamic pricing system can help sellers remain competitive, manage inventory, respond to demand fluctuations, and improve revenue. Existing research has investigated regression models, ensemble learning, reinforcement learning, and other machine-learning approaches for predicting or optimizing prices. Research on dynamic pricing has also explored reinforcement-learning and game-theoretic approaches for e-commerce platforms [4]. These studies demonstrate that automated pricing can respond to complex market conditions more effectively than fixed pricing strategies.



However, dynamic pricing also introduces an important concern: pricing transparency. A product may appear to be offered at a substantial discount even when the displayed original price does not accurately represent the product's normal historical price. For example, a seller may display a high strike-through price and advertise a large percentage discount even though the product may have been sold close to the current price for most of the preceding period. Such practices can make consumers believe that they are receiving a larger benefit than they actually are.

This problem is particularly relevant during large online shopping events, festival sales, seasonal promotions, and flash-sale campaigns. Consumers commonly compare the current selling price with the displayed original price instead of examining the complete historical price behaviour of the product. Therefore, simply calculating the percentage difference between the current selling price and the displayed original price is not sufficient to determine whether a discount is genuine. A more reliable approach requires historical price analysis and a comparison between the claimed reference price and the product's normal pricing range.

Existing machine-learning systems have addressed related problems from different perspectives. Some systems focus primarily on predicting fair prices, while others concentrate on identifying anomalous listings or comparing prices across multiple online marketplaces. The study Dynamic Pricing for Anomaly Detection in Online Market Places investigates fraud detection by predicting listing prices and identifying discrepancies between predicted and listed prices [1]. Similarly, IntelliPrice reviews the combination of artificial intelligence and web-scraping techniques for dynamic e-commerce price prediction and comparison [2]. These studies demonstrate the usefulness of machine learning and automated data collection, but they do not directly provide a unified mechanism for validating whether a displayed discount is genuine using historical price behaviour.

Another important research direction is seasonality-aware pricing. Product demand and prices can vary significantly during holidays, festivals, weekends, promotional periods, and other seasonal events. Recent work on a Seasonality Pricing Recommendation Engine combines competitor web scraping, catalogue normalization, seasonality analysis, and price optimization to support retail pricing decisions [3]. Such approaches demonstrate that temporal and competitive information can improve pricing recommendations, but the question of whether a promotional discount is honestly represented to the consumer remains a separate challenge.

Machine learning in e-commerce is increasingly being applied to recommendation systems, fraud detection, demand forecasting, customer segmentation, personalization, and dynamic pricing. A recent IEEE Access survey reviews supervised, unsupervised, reinforcement, and hybrid learning approaches in e-commerce and identifies challenges including explainability, scalability, and real-time adaptability [5]. These challenges are particularly relevant to automated pricing systems because pricing decisions can directly influence both sellers and consumers.

To address these limitations, this paper proposes Fairlytics: AI-Based Intelligent Dynamic Pricing With Discount Authenticity Validation. Fairlytics combines two closely related functions within a unified architecture: a Dynamic Pricing Engine (DPE) and a Discount Authenticity Validator (DAV). The DPE estimates a fair market price using historical sales information, competitor prices, inventory levels, product characteristics, and seasonal signals. The DAV then evaluates the authenticity of an advertised discount by comparing the claimed original price with historical price behaviour and the independently estimated fair price.

The proposed system uses an ensemble-based pricing approach together with anomaly detection and rule-based validation. The ensemble pricing model combines multiple machine-learning techniques to improve robustness, while the Discount Authenticity Validator uses historical price information and an Isolation Forest-based anomaly detection mechanism to identify unusual price patterns. The outputs of these components are combined to generate an explainable Fairness Score that provides evidence for the resulting decision.

The objective of Fairlytics is not simply to identify whether a price is high or low. Instead, the system attempts to determine whether the displayed price and discount representation are consistent with historical market behaviour. This distinction is important because a discount can appear mathematically large while providing little actual price benefit to the consumer. By reconstructing a product's price trajectory and comparing it with the claimed reference price, Fairlytics aims to provide a more meaningful measure of discount authenticity.

Fairlytics is designed to support three primary groups of users. Consumers can use transparency information to better understand whether a displayed discount appears genuine. Sellers can use fair-price predictions to make more informed pricing decisions. Platform administrators can use the system to identify suspicious listings and review the evidence



behind discount-validation decisions. This multi-user design extends the usefulness of pricing analytics beyond revenue optimization alone.

The remainder of this paper is organized as follows. Section II presents the background concepts related to dynamic pricing, ensemble learning, reference-price manipulation, anomaly detection, and explainable fairness scoring. Section III discusses related research in dynamic pricing, web scraping, seasonality-aware pricing, and machine learning in e-commerce. Section IV presents the comparative analysis and research gap. Section V describes the proposed Fairlytics architecture. Section VI explains the methodology and algorithms used in the system. Section VII describes the prototype implementation. Section VIII presents the results and evaluation approach. Section IX presents the discussion and practical implications. Section X describes the advantages of the proposed system. Section XI presents the conclusion, and Section XII describes the future scope of the proposed work.

II. BACKGROUND CONCEPTS

A. Dynamic Pricing and Ensemble Learning

Dynamic pricing is the practice of continuously adjusting a product's price in response to demand, competition, inventory, and time-based signals rather than holding a fixed catalog price. Ensemble learning combines several base models — such as Linear Regression, Random Forest, Gradient Boosting, and neural networks — through techniques like stacking, so that the combined prediction is typically more accurate and more robust to category-specific noise than any single model alone.

B. Discount Framing and Reference-Price Manipulation

A discount is normally displayed as a percentage or amount off a stated “original” or “strike-through” reference price. Reference-price manipulation occurs when a seller inflates this original price shortly before a sale so that the advertised discount looks larger than the real change in price. Detecting this requires comparing the claimed original price against the product's genuine historical price trajectory, not just against the current listed price.

C. Anomaly Detection for Price Histories

Anomaly detection identifies data points that deviate significantly from an expected pattern. Applied to price histories, an unsupervised model such as Isolation Forest can learn the normal range of a product's price over time and flag a claimed original price that falls far outside that range, without requiring a large hand-labeled set of known fraud cases.

D. Explainable Fairness Scoring

Explainability, in this context, means that every automated pricing or fraud decision is accompanied by the concrete evidence that produced it — for example, a historical price chart or a specific deviation value — rather than a single opaque score. This is essential for a system, like Fairlytics, whose output is reviewed by human administrators and shown directly to consumers.

III. LITERATURE REVIEW

The reviewed literature is grouped into four themes: (A) dynamic pricing and price-anomaly fraud detection, (B) live web-scraping and price-comparison systems, (C) seasonality-aware demand forecasting, and (D) broader machine-learning-in-e-commerce surveys.

A. Dynamic Pricing and Price-Anomaly Fraud Detection

Dynamic pricing research has traditionally focused on maximizing seller revenue. Rule-based and reinforcement-learning approaches adjust prices in response to inventory levels, competitor pricing, and time-to-expiry of perishable or seasonal stock [1]. More recent systems incorporate ensemble regression techniques, combining Random Forest, Gradient Boosting, and Decision Tree models to predict a fair product price from structured attributes such as category, brand, and historical sales trends, and treat a large deviation between the predicted and listed price as a fraud signal [2]. This line of work targets listing-level price manipulation broadly, rather than the specific and increasingly common tactic of inflating a reference price to manufacture a discount.



B. Live Web-Scraping and Price-Comparison Systems

A parallel line of work uses web scraping and concurrent data collection to build real-time price-comparison engines. These systems retrieve live listings from multiple marketplaces, using tools such as BeautifulSoup and Selenium behind a lightweight backend, and surface the lowest available price to the consumer [3]. Such systems demonstrate the feasibility of continuously refreshed pricing data, but they generally stop at comparison and do not evaluate whether an individual seller's discount claim is truthful.

C. Seasonality-Aware Demand Forecasting

Seasonality-aware pricing engines show that incorporating holiday calendars, weather data, and promotional history into demand forecasting models — such as Random Forest, XGBoost, and LSTM networks — improves both revenue and inventory efficiency [4]. These systems optimize the price level itself but are not designed to audit how that price is subsequently framed as a discount to the consumer.

D. Machine Learning in E-Commerce: Broader Surveys

Broader surveys of machine learning in e-commerce confirm that dynamic pricing, fraud detection, and personalization are converging research areas, and identify explainability, real-time adaptability, and integrated multi-task frameworks as open challenges [5]. Collectively, this body of work establishes strong individual foundations in demand-responsive pricing, live data collection, and price-anomaly-based fraud detection. What remains largely unaddressed is a system that treats discount authenticity as a first-class prediction target alongside price itself, using a shared, auditable model of a product's genuine price history.

IV. COMPARATIVE ANALYSIS AND RESEARCH GAP

Table I compares the reviewed systems against Fairlytics along five dimensions: core focus, pricing approach, how (if at all) discount or fraud claims are handled, whether outputs are explainable, and whether a consumer-facing transparency signal is produced. No reviewed system addresses more than two of these dimensions at once, which motivates a single, unified design.

TABLE I. COMPARATIVE ANALYSIS OF REVIEWED SYSTEMS

Study	Core Focus	Pricing Approach	Discount / Fraud Handling	Explainable Output	Consumer-Facing
Dynamic Pricing Anomaly Detection [1]	Fraud detection via price deviation	Regression-based fair-price prediction	Flags large predicted-vs-listed price gaps	Partial	No
Ensemble Price Prediction [2]	Fair-price prediction	Random Forest / Gradient Boosting ensemble	Not addressed	Partial	No
IntelliPrice / Live Scraping [3]	Real-time price comparison	Live scraped price aggregation	Not addressed	No	Yes (price only)
Seasonality Pricing Engine [4]	Demand-aware price optimization	Random Forest / XGBoost / LSTM	Not addressed	No	No
ML in E-Commerce Survey [5]	Cross-cutting ML applications	Reviews multiple paradigms	Discussed conceptually only	Discussed conceptually only	No
Fairlytics (Proposed)	Pricing + discount authenticity	Stacked ensemble (DPE)	Price-history anomaly + rule-based DAV	Yes — evidence-	Yes — transparency badge



Study	Core Focus	Pricing Approach	Discount / Fraud Handling	Explainable Output	Consumer-Facing
				backed Fairness Score	

The comparison in Table I highlights five recurring gaps that Fairlytics is designed to close:

- Dynamic pricing systems generally optimize revenue without validating whether the resulting price is presented to consumers through a truthful discount frame.
- Existing fraud detection work flags anomalies between predicted and listed prices but rarely models the historical price trajectory needed to catch inflated “strike-through” reference prices.
- Live price-comparison and web-scraping systems refresh price data continuously but stop at comparison, without judging whether an individual seller's discount claim is truthful.
- Discount legitimacy checks, where they exist, tend to be manual or rule-based, and do not scale to the volume of promotions run on modern marketplaces.
- Few systems expose a consumer-facing transparency signal; fraud detection outputs are typically used only internally by platform operators.

Fairlytics is motivated by closing these gaps through a single architecture that shares data, features, and a common notion of “fair price” across both the pricing and discount-validation tasks.

V. PROPOSED SYSTEM: FAIRLYTICS ARCHITECTURE

A. System Overview

Fairlytics is organized as a five-stage pipeline: data ingestion, feature engineering, a dual prediction core, a fairness-scoring layer, and an administrator/consumer-facing dashboard. The dual prediction core contains two cooperating engines that share the same feature store — a Dynamic Pricing Engine (DPE) and a Discount Authenticity Validator (DAV) — so that the fair-price estimate produced by the DPE becomes the reference point the DAV uses to judge whether an advertised discount is genuine. Fig. 1 illustrates the end-to-end pipeline.

Data Collection (sales, prices, inventory, promotions)
Preprocessing & Feature Engineering
Dynamic Pricing Engine (ensemble regression)
Discount Authenticity Validator (price-history + anomaly model)
Fairness Scoring & Explanation Layer
Admin / Seller / Consumer Dashboard

Fig. 1. Fairlytics five-stage pipeline architecture.

B. Dynamic Pricing Engine (DPE)

The DPE predicts a fair market price for each product from structured features (category, brand, inventory level, competitor prices) and time-based signals (day of week, season, upcoming holidays, sales velocity). It uses a stacked ensemble of Linear Regression, Random Forest, Gradient Boosting, and a lightweight neural network, with model stacking combining their outputs into a final prediction. The DPE recomputes prices on a scheduled interval and whenever a significant change in demand or inventory is detected.

C. Discount Authenticity Validator (DAV)

The DAV addresses the core novelty of Fairlytics. For each promoted listing, it reconstructs the product's historical price series and compares three quantities: the seller's claimed “original” price, the DPE's independently predicted fair price, and the statistically observed price over a trailing window (for example, the preceding 90 days). A discount is scored as authentic when the claimed original price is consistent with the observed historical price band; it is flagged



when the claimed original price sits well above both the historical trend and the model-predicted fair price. The DAV uses a time-series anomaly detection model (an Isolation Forest applied to price-history features) combined with a transparent rule layer, so that every flag can be explained in terms of concrete, auditable evidence rather than an opaque score alone.

D. Fairness Scoring and Dashboard

Outputs from the DPE and DAV are combined into a single Fairness Score per listing, ranging from a fully authentic, well-priced item to a flagged listing requiring seller review. Administrators see a ranked queue of suspicious listings with the specific evidence (historical price chart, predicted price, claimed discount) driving each flag. Consumers see a simplified transparency badge on the product page.

VI. METHODOLOGY

The methodology follows a standard supervised and semi-supervised learning workflow, adapted to the dual-engine design.

A. Data Collection and Preprocessing

Historical sales records, listed prices, claimed discount percentages, inventory snapshots, and promotional calendars are collected and merged into a unified dataset. Missing values are imputed, duplicate listings are removed, and prices are normalized per product category to control for scale differences across categories.

B. Feature Engineering

Feature sets include: (i) structured attributes such as category, brand, and rating; (ii) temporal features such as day-of-week, proximity to holidays, and sales-event flags; (iii) price-history features such as rolling mean, rolling standard deviation, and the maximum price observed over trailing 30/60/90-day windows; and (iv) discount features such as claimed discount percentage and the ratio between claimed original price and the trailing price statistics.

C. Model Training

The DPE is trained as a regression task with mean absolute percentage error (MAPE) as the primary loss metric, using k-fold cross-validation across product categories to avoid category-specific overfitting. The DAV is trained as a hybrid task: an unsupervised Isolation Forest identifies price-history anomalies without requiring labeled fraud data, while a supervised classifier — trained on a smaller set of manually reviewed authentic and inflated discount examples — refines the final authenticity decision.

D. Algorithm Summary

Algorithm 1 summarizes the end-to-end scoring procedure applied to each listing at inference time:

Algorithm 1: Fairlytics Listing Evaluation

Input: product listing L with claimed original price P_o and claimed discount d

1. Extract feature vector x from L using the feature-engineering pipeline.
2. Compute fair price estimate $P_f = \text{DPE}(x)$.
3. Retrieve trailing price history H for the product.
4. Compute historical price band $[H_{\min}, H_{\max}]$ from H .
5. Compute anomaly score $a = \text{IsolationForest}(H, P_o)$.
6. If P_o is within a tolerance band of $H_{\max}$ and P_f , mark discount as Authentic.
7. Else if a exceeds the anomaly threshold, mark discount as Flagged and attach supporting evidence.
8. Else mark discount as Under Review for manual/admin verification.
9. Output Fairness Score combining P_f , a , and the discount decision.



VII. IMPLEMENTATION DETAILS

A prototype implementation of Fairlytics is designed around a Python-based machine learning backend and a lightweight web stack for the dashboard. Scikit-learn provides the regression ensemble and Isolation Forest implementations; pandas and NumPy handle data preprocessing; a Flask or Django REST API exposes pricing and authenticity endpoints; and a React-based frontend renders the administrator dashboard and the consumer-facing transparency badge. Historical and transactional data is stored in a relational database (MySQL/PostgreSQL) for structured records, with a document store (MongoDB) optionally used for raw scraped listing data prior to preprocessing. Scheduled batch jobs (e.g., via Celery or a cron-based scheduler) periodically refresh price predictions and re-score active promotions.

VIII. RESULTS AND EVALUATION

Fairlytics is currently presented as a prototype-stage system design; therefore, the evaluation focuses on measurable criteria for validating both the pricing engine and the discount-authenticity module. The evaluation is designed to determine whether the proposed architecture can accurately estimate fair prices, identify suspicious discount representations, and provide understandable evidence to users and administrators.

The first evaluation component is the performance of the Dynamic Pricing Engine (DPE). Pricing accuracy will be measured by comparing the predicted fair price with the actual observed transaction price. Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Mean Absolute Percentage Error (MAPE) can be used to evaluate prediction performance. Lower error values indicate that the predicted fair price is closer to the observed market price. Evaluation will be performed across different product categories to determine whether the model remains reliable when product characteristics and price ranges vary.

The second evaluation component focuses on the Discount Authenticity Validator (DAV). The DAV will be evaluated using a dataset containing examples of authentic discounts and potentially inflated reference-price discounts. The system's predictions will be compared with manually reviewed labels. Precision, recall, F1-score, and accuracy will be used to measure classification performance. Precision is important because incorrectly flagging legitimate discounts could negatively affect sellers, while recall is important because failing to identify manipulated discounts could reduce consumer protection.

The third evaluation component is the historical price anomaly detection mechanism. For each product, historical price observations will be analysed over a defined time window. The system will identify unusually high or low values relative to the normal historical range. The Isolation Forest component will generate an anomaly score, which will then be combined with rule-based checks involving the claimed original price and predicted fair price. This hybrid approach is intended to reduce dependence on a single threshold and provide more interpretable decisions.

The fourth component evaluates the Fairness Score and explanation layer. Instead of displaying only a binary result such as genuine or flagged, Fairlytics provides supporting evidence for its decision. This evidence can include the predicted fair price, historical price range, claimed original price, current selling price, discount percentage, and anomaly score. The explanation layer can therefore be evaluated based on whether an administrator can understand the reason for a flagged listing without inspecting the underlying machine-learning model.

The system will also be evaluated under different market conditions. These conditions include stable prices, rapidly changing prices, seasonal demand changes, promotional periods, inventory shortages, and competitor price fluctuations. Such testing is important because an anomaly detector that performs well under stable prices may incorrectly flag legitimate price changes during seasonal sales or high-demand periods.

A separate evaluation will examine the effect of the historical time window used by the DAV. For example, 30-day, 60-day, and 90-day historical windows can be compared. A shorter window may respond more quickly to recent market changes, while a longer window can provide a more stable representation of normal pricing behaviour. The optimal window can therefore be selected based on validation performance.

Dashboard usability will also be evaluated because the proposed system is intended for practical use by administrators and consumers. Administrators should be able to identify suspicious listings, view supporting historical-price evidence, and determine whether manual review is necessary. Consumers should be able to understand the transparency badge



without requiring technical knowledge of machine learning. Task-completion time and user feedback can be used as practical usability measures.

The evaluation will finally compare the proposed hybrid approach with simpler baseline methods, such as fixed-threshold discount detection, regression-only price prediction, and anomaly detection without the rule-based validation layer. This comparison will help determine whether combining multiple components provides a measurable improvement over isolated approaches.

The expected outcome is that Fairlytics will provide a more comprehensive approach to pricing transparency by combining fair-price prediction and discount authenticity validation. However, quantitative performance claims should only be reported after testing the implemented system on a suitable dataset. Since the current work is at the prototype/design stage, final numerical results will be presented after data collection, model training, validation, and experimental testing are completed.

F. Prototype Interface

To demonstrate feasibility, a working prototype of the Fairlytics interface was implemented. Fig. 2 shows the landing page, where the system reports a Fairness Score alongside a comparison of listed prices across platforms, distinguishing a genuine best deal from listings with suspicious or inflated discounts. Fig. 3 shows the marketplace-coverage view, which aggregates listings from multiple partner platforms for comparison. Fig. 4 shows the product-level comparison view, which highlights the recommended genuine best deal together with the computed savings. Fig. 5 shows the full listings table for a single product across four platforms, with the authenticity verdict and supporting analysis generated for each listing.




Fig. 2. Fairlytics landing page showing the AI-computed Fairness Score and a genuine-vs-suspicious price comparison across platforms.



Fig. 3. Marketplace-coverage view aggregating listings from partner platforms for price comparison.

**Find the Best Deal**

Select a product and compare prices across marketplaces with real-time discount validation.

Choose Product from Verified Catalog

Cetaphil Gentle Foaming Daily Facial Cleanser

Compare Prices

RECOMMENDED GENUINE BEST DEAL

SKINCARE

Cetaphil Gentle Foaming Daily Facial Cleanser

Best available on **Amazon**

₹499.00

MRP: ₹575.00 **Save ₹76.00 (13.2% OFF)**

✓ GENUINE DEAL

Fig. 4. Product-level comparison view highlighting the recommended genuine best deal and computed savings.

All Marketplace Listings (4 Platforms) Sorted by selling price (Lowest to Highest)

PLATFORM	SELLING PRICE	LISTED MRP	DISCOUNT	MRP MULTIPLE	AUTHENTICITY VERDICT	ANALYSIS DETAILS
AmazonBest Pick	₹499.00	₹575.00	13.2%	1.15x	✓ GENUINE	Authentic pricing: reasonable discount (13.2%) and realistic MRP multiple (1.15x).
Purple	₹510.00	₹575.00	11.3%	1.13x	✓ GENUINE	Authentic pricing: reasonable discount (11.3%) and realistic MRP multiple (1.13x).
Nykaa	₹515.00	₹575.00	10.4%	1.12x	✓ GENUINE	Authentic pricing: reasonable discount (10.4%) and realistic MRP multiple (1.12x).
NewU	₹545.00	₹575.00	5.2%	1.06x	✓ GENUINE	Authentic pricing: reasonable discount (5.2%) and realistic MRP multiple (1.06x).

Fig. 5. Full marketplace-listings table with per-platform authenticity verdict and supporting analysis.

IX. DISCUSSION

Two practical takeaways stand out from this design process. First, the underlying components — ensemble regression for price prediction and Isolation Forest for anomaly detection — are already well-established and widely available in libraries such as scikit-learn, so Fairlytics does not need to invent new algorithms; its contribution is in how these components are combined around a shared notion of fair price rather than in any single new model. Second, the choice of how strict the Discount Authenticity Validator should be involves a genuine trade-off rather than a default setting: a low anomaly threshold catches more manipulated discounts but risks flagging legitimate clearance sales or genuinely volatile-price items, while a high threshold reduces false positives at the cost of missing subtler manipulation. As Table I shows, systems that emphasize strict fraud detection (comparable to [1]) and systems that emphasize smooth, uninterrupted pricing (comparable to [2], [4]) sit at different points on this trade-off; Fairlytics addresses it by routing borderline cases to an “Under Review” state for human administrator judgment rather than forcing a fully automated binary decision, echoing the general observation from the broader e-commerce ML literature that human-in-the-loop review remains important wherever automated decisions directly affect consumer trust [5].

X. ADVANTAGES OF THE PROPOSED SYSTEM

- Unifies dynamic pricing and discount-fraud detection in a single, shared-feature architecture instead of two disconnected systems.
- Produces an explainable Fairness Score backed by concrete historical-price evidence, rather than an opaque anomaly flag alone.
- Gives consumers a direct transparency signal, extending fraud detection benefits beyond internal platform operations.
- Modular design allows the DPE and DAV to be retrained or extended independently as new categories or markets are added.



- Hybrid unsupervised/supervised training for discount validation reduces dependence on large volumes of manually labeled fraud data.

XI. CONCLUSION

Fairlytics presents an AI-based framework for combining intelligent dynamic pricing with discount authenticity validation in e-commerce environments. The proposed system addresses two closely connected challenges: estimating a fair and demand-responsive product price and determining whether an advertised discount is consistent with the product's historical pricing behaviour. Instead of treating these problems as independent tasks, Fairlytics integrates them through a shared data and feature pipeline.

The proposed architecture consists of a Dynamic Pricing Engine, a Discount Authenticity Validator, a Fairness Scoring and Explanation Layer, and dashboards for administrators, sellers, and consumers. The Dynamic Pricing Engine uses product, market, inventory, competitor, and temporal information to estimate a fair price. The Discount Authenticity Validator analyses historical price trajectories and uses anomaly detection together with transparent rules to identify potentially misleading reference prices.

A major advantage of the proposed approach is explainability. Rather than producing only a binary fraud decision, Fairlytics can provide supporting evidence such as historical price ranges, predicted fair prices, claimed original prices, and anomaly scores. This information can assist administrators in reviewing suspicious listings and can provide consumers with a clearer understanding of advertised discounts.

The proposed framework is currently positioned as a prototype-stage system, and its final effectiveness depends on evaluation using real or appropriately constructed marketplace data. Future experimental validation should measure pricing accuracy, discount-detection precision and recall, anomaly-detection performance, and dashboard usability. Overall, Fairlytics provides a unified foundation for improving pricing transparency, supporting data-driven pricing decisions, and increasing consumer confidence in e-commerce discount claims.

XII. FUTURE SCOPE

The proposed Fairlytics framework can be extended in several directions. First, the system can be evaluated using larger and more diverse real-world datasets covering multiple e-commerce categories, sellers, geographical markets, and pricing conditions. A larger dataset would allow the models to learn more representative pricing patterns and improve generalization across different product categories.

Second, the Discount Authenticity Validator can be enhanced using additional signals beyond numerical price history. Product images, product descriptions, customer reviews, seller ratings, promotional text, and review sentiment could be incorporated to identify inconsistencies between promotional claims and actual product information.

Third, the Dynamic Pricing Engine can be extended using reinforcement learning. Reinforcement-learning-based pricing has been investigated for e-commerce environments, including approaches that learn pricing strategies from market interaction [4]. Such techniques could allow Fairlytics to adapt pricing recommendations continuously as demand, competition, and inventory conditions change.

Fourth, future versions can incorporate real-time competitor-price monitoring through responsible web-data collection and APIs where available. This would allow the pricing engine to react more quickly to changes in marketplace conditions. The integration of seasonality information, competitor prices, and automated data collection is also consistent with recent research in retail pricing systems [3].

Fifth, the fairness mechanism can be further developed to consider consumer-oriented constraints. The system could evaluate whether price changes remain within predefined transparency or fairness limits and could provide alerts when pricing behaviour becomes unusually volatile.

Finally, a user study can be conducted to determine whether transparency badges and historical-price explanations actually improve consumer understanding and purchasing decisions. Such studies would help evaluate Fairlytics not only from a machine-learning perspective but also from a usability and consumer-trust perspective.



ACKNOWLEDGMENT

The authors thank Rashmi and Lavanya N L for their guidance and continuous support throughout this project, and the Department of Computer Science and Engineering, East West College of Engineering, for providing the resources necessary to carry out this work.

DECLARATIONS

- **Funding Source:** No funding was received for this work.
- **Acknowledgement:** The authors thank East West College of Engineering and the Department of Computer Science and Engineering for their support.
- **Data Availability:** Data is available from public sources or upon reasonable request.
- **Authors' Contribution:** All authors contributed to the research, implementation, analysis, and manuscript preparation.
- **Use of AI and AI-Assisted Technologies:** AI tools were used for language improvement and grammar checking. The authors verified the content.
- **Conflict of Interest:** The authors declare no conflict of interest.
- **Copyright Permissions:** Copyright permissions were obtained for all figures and tables used.

REFERENCES

- [1] C. Krishnapriya, G. P. Sajeev, T. R., S. K. B., L. E. Sunny, and J. A. Mathew, "Dynamic Pricing for Anomaly Detection in Online Market Places," *International Journal for Research in Applied Science and Engineering Technology*, vol. 12, no. 5, pp. 1323–1329, 2024, doi: 10.22214/ijraset.2024.61754.
- [2] M. Senbagavalli, N. Kumar, A. Sungheetha, D. Kumar, and S. Ghosh, "IntelliPrice: A Unified Review of AI and Web Scraping Techniques for Dynamic E-Commerce Price Prediction and Comparison," in *Proc. 2025 IEEE 6th India Council International Subsections Conf. (INDISCON)*, 2025, doi: 10.1109/INDISCON66021.2025.11252154.
- [3] K. Govardhanan, "Seasonality Pricing Recommendation Engine for Retail: A Price Match System Using Competitor Web Scraping and Optimized Pricing Suggestions," in *Proc. 2025 4th Int. Conf. on Applied Artificial Intelligence and Computing (ICAAIC)*, 2025, doi: 10.1109/ICAAIC64647.2025.11331171.
- [4] C. Yin and J. Han, "Dynamic Pricing Model of E-Commerce Platforms Based on Deep Reinforcement Learning," *Computer Modeling in Engineering & Sciences*, vol. 127, no. 1, pp. 291–307, 2021, doi: 10.32604/cmes.2021.014347.
- [5] E. Dritsas and M. Trigka, "Machine Learning in E-Commerce: Trends, Applications, and Future Challenges," *IEEE Access*, vol. 13, pp. 99048–99067, 2025, doi: 10.1109/ACCESS.2025.3572865.
- [6] F. E. Isinkaye, Y. O. Folajimi, and B. A. Ojokoh, "Recommendation Systems: Principles, Methods and Evaluation," *Egyptian Informatics Journal*, vol. 16, no. 3, pp. 261–273, 2015.
- [7] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation Forest," in *Proc. 8th IEEE Int. Conf. on Data Mining*, 2008, pp. 413–422.
- [8] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," in *Proc. 22nd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, 2016, pp. 785–794.
- [9] A. Kastius and R. Schlosser, "Dynamic Pricing under Competition Using Reinforcement Learning," *Journal of Revenue and Pricing Management*, vol. 21, pp. 50–63, 2022, doi: 10.1057/s41272-021-00285-3.
- [10] Y. Liu, T. Cui, G. Li, T. R. Payne, Y. Yue, and K. L. Man, "A Multimodal Deep Reinforcement Learning Framework for Dynamic Pricing Optimisation in E-Commerce," *IEEE Access*, 2026, doi: 10.1109/ACCESS.2026.3680140.