Data-Driven Pricing Strategy in E-Commerce: Predictive Modeling Using Consumer Purchase Behavior

Rangasamy S
PG Student
Department of Management Studies
K S Rangasamy College of Technology
imrangasamyrs@gmail.com

Lavanya RG
PG Student
Department of Management Studies
K S Rangasamy College of Technology
lkglavanya17@gmail.com

Yogeswaran C
PG Student
Department of Management Studies
K S Rangasamy College of Technology
yogiwaran827@gmail.com

Hemamalini R
Professor
Department of Management Studies
K S Rangasamy College of Technology
hemamalini@ksrct.ac.in

***Abstract*— Adaptive, data-driven pricing is a necessity for online merchants experiencing volatile demand and intense competition. This research suggests an integrated predictive model that uses historical transaction logs, click-stream traces, and exogenous market indicators to infer best price points at SKU, segment, and session levels. The pipeline in this model begins with Recency-Frequency-Monetary (RFM) features and K-Means/Hierarchical clustering to obtain behaviorally meaningful customer segments; then elastic-net regression predicts short-term price-elasticity for each segment; lastly, a reinforcement-learning layer adjusts prices in near-real-time to maximize expected contribution margin subject to inventory and competitor-price constraints. We test the framework on 2.8 million orders for a mid-size fashion e-retailer during January 2023 – December 2024. Our system compares favorably with the company's rule-based approach. It increases gross profit by 7.9 %, conversion among high-lifetime-value customers by 5.4 %, and reduces markdown expenditure by 11.2 %. Robustness checks under severe demand shocks—e.g., flash sales and influencer-driven traffic spikes—verify steady performance. The contribution is two-fold: (i) methodological—integrating segmentation, econometric elasticity, and machine-learning control within one loop; (ii) managerial—showing how granular behavioral data can translate to defensible margin gains with customer goodwill intact. Ethical and regulatory aspects of personalized prices are also examined**

Keywords— Dynamic pricing, E-commerce analytics, Customer segmentation, Recency-Frequency-Monetary (RFM), Price elasticity, Elastic-net regression, Reinforcement learning, Predictive modelling, Revenue management, Behavioral data, Machine-learning control, Personalized pricing.

I. INTRODUCTION (*HEADING 1*)

Price is the most direct control an e-commerce company can manipulate, but it is also the hardest to fine-tune continuously in scale. Historical markdown calendars, coupon ladders, or A/B tests deal with customers as if they were homogenous and do not account for quickly changing circumstances like social-media buzz or competitor scraping bots. Innovations in cloud computing, low-latency data pipes, and machine-learning (ML) algorithms allow retailers to process millions of micro-decisions per minute and customize offers at the level of a single cart. Dynamic pricing, if implemented responsibly, thus holds out promise for concurrent increase in consumer surplus and enterprise margin. The practice is contentious, though: regulators now question whether ubiquitous data gathering will result in discriminatory or exploitative prices

WS This tension drives the current research.
We locate our research at the intersection of revenue management and customer analytics. Drawing on evidence that various buyer archetypes have varying elasticities, we introduce a two-stage approach: segment first, price across segments. Clustering separates groups that are heterogenous in willingness-to-pay, search effort, and responsiveness to delay in delivery. In each segment we place predictive regressions that render contextual covariates—device, channel of referral, and day-of-week, among others—onto conversion probability and basket value as a function of price. Prices are chosen through an optimization routine that maximizes expected profit within fairness and inventory constraints.

Our contributions are threefold. First, we combine traditional econometric elasticity estimation with contemporary ML and reinforcement learning in an end-to-end system deployable on commodity cloud infrastructure. Second, we present a rigorous empirical assessment on a two-year transactional dataset, thus capturing pandemic-induced distortions and seasonal peaks. Third, we demonstrate that the same analytical backbone can serve adjacent fields—above all, aviation safety—by adding, for cross-validation purposes, a comparative bar chart of model accuracy on a hard-landing detection task (Figure 3), highlighting the portability of ML evaluation methods across verticals.

II. LITERATURE SURVEY

Dynamic pricing has evolved markedly over the past two decades, moving from rule-based heuristics to data-intensive optimisation that exploits advances in machine learning (ML) and cloud computing. The earliest scholarship framed dynamic pricing largely as an extension of electronic catalogues. Kannan and Kopalle argued that the Internet lowers menu costs and allows near-frictionless price revisions, but warned that excessive variability could erode consumer trust [3]. Their behavioural insights still undergird much of the algorithmic-pricing debate.

2.1 Foundations in customer analytics

Reliable, individual-level demand estimation depends on good customer analytics. Recency-Frequency-Monetary (RFM) models are still a workhorse because they compress longitudinal behaviour into a compact summary. Verhoef and Donkers showed that integrating RFM with acquisition models enhances lifetime-value targeting [4]. Fader, Hardie, and Lee codified RFM using iso-value curves, offering a geometric representation which enables segment-level strategy [5]. Zhang et al. more recently used deep embeddings to represent high-dimensional purchase histories for personalized pricing, with statistically significant lift above traditional RFM on a large e-retail dataset [15].

2.2 Demand estimation and price elasticity

Personalized Accurate elasticities are critical due to the fact that minor prediction errors translate into major profit fluctuations. Montgomery made the case for elasticities to be understood from a behavioural perspective to prevent the "illusion of precision" afflicting purely statistical models [6]. Elastic-net regression provides a principled approach to addressing multicollinearity in high-dimensional spaces; Bertschek et al. demonstrated that it performs better than LASSO and Ridge on German online-grocery data [7]. Gradient-Boosting Machines (GBM) perform well on mixed categorical–numeric inputs too; Guo and Wang achieved a 12 % RMSE improvement over Random Forests in demand forecasting for a Chinese electronics retailer [8]. Park and Gupta warn, though, that time-series cross-validation is necessary; naïve random splits will inflate apparent accuracy by up to 30 % [17].

2.3 Dynamic pricing with competition

Revenue-management theory provides the normative foundation for pricing. Talluri and Van Ryzin's monograph integrated control-theoretic methods with stock constraints [9], while Bitran and Caldentey enumerated industry case studies in price models ranging from airlines to hotels [10]. Besbes et al. developed demand-learning mechanisms that dynamically adjust prices in the presence of parameter uncertainty [11]. Empirical evidence substantiates these theories: Bapna et al. chronicled how emerging-market airlines differentiate fares by channels in balancing price stimulus of price-sensitive segments with safeguarding premium yield [14].

Competition makes the challenge more severe. Kastius and Schlosser simulated duopolistic price wars using reinforcement learning (RL), identifying non-stationary equilibria in which exploratory actions occasionally shatter tacit collusion [2]. Oh and Dekker applied RL to an actual-time grocery environment, incorporating perishability constraints and showing 6 % profit gain [25]. A second line of study investigates marketplace platforms: Chen, Mislove, and Wilson identified algorithmic repricers on Amazon to converge to leader–follower patterns, posing questions about unintentional price coordination [12].

2.4 Fairness and transparency of algorithms

The increasing prevalence of black-box models in high-stakes decisions has prompted demands for interpretability. Rudin contends that intrinsically interpretable models can substitute post-hoc accounts wherever possible [18]. Fairness constraints are coming to pricing research: Mehta and Singh offer a convex optimisation framework capping price differences between protected groups, with little revenue loss in simulations [16]. Wachter and Mittelstadt outline an ethical framework based on EU consumer law principles of transparency, contestability, and proportionality [19]. Policy-wise, Cowgill distils evidence that algorithmic pricing can positively and negatively impact consumer welfare depending on market structure [13], while Narayanan warns antitrust tools are potentially not equipped to spot algorithmic collusion [24].

2.5 Robustness, privacy, and covariate shift

Precision is driven by fine-grained data but increases risk to privacy. Abadi et al. present how differential-privacy methods can protect sensitive buying records with very little loss of accuracy when $\varepsilon \leqslant 1$ [23]. Jellinek illustrates that covariate-shift detection is essential in commercial data; undetected distributional drift lowered a high-performing price-recommendation system's profit by 15 % within six months [22].

2.6 Synthesis and research gap

But, Together, the literature coalesces around three conclusions: Segmentation is key—ranging from traditional RFM to embeddings, clustering buyers by behaviour always increases elasticity estimation.RL beats static rules in non-stationary competitive settings but needs guardrails to avoid unfair results.Transparency and privacy concerns are now first-order constraints that influence possible algorithm design.Despite that, few works combine segmentation, elasticity regression, and RL under a single closed loop while ensuring stringent auditing of fairness and privacy. This research bridges this gap by implementing an end-to-end framework that is tested on a two-year e-retail dataset, showing new evidence that profit boosts can go hand-in-hand with ethical protection

III. Research Methodology

This section describes the architecture and implementation of the predictive price framework. It uses a multi-stage pipeline that combines behavioral segmentation, demand estimation, and reinforcement learning-based price optimization. The entire pipeline is depicted in Figure 4 – Methodology Pipeline

3.1 Sources of Data

We collaborated with Fashion Hub, an average-sized European online fashion retailer, and pulled 24 months of data for:

Transactions: sales at SKU level with timestamp, list price, discount, coupon, gross margin, and on-hand stock.

Behavioural logs: page views, add-to-cart, dwell time, searched terms, referral channel, device, and geo-granularity at NUTS-2 region.

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 12 | Dec - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

Exogenous signals: competitor prices through daily scraping, Google Trends indexes, and macro indicators like consumer confidence.

Following GDPR-compatible anonymisation and filtering out of outliers (Cook's D > 4/n), the final corpus contained 2 807 346 orders and 182 million click events. A. Unified Customer Data Platform (CDP)

3.2 Feature Engineering

We designed three families of features: Customer state vectors (per session): RFM scores, loyalty-tier, coupon history, and return ratio.

Contextual covariates: time-of-day, day-of-week, special events (Black Friday flag), remaining size range, and competitor price gap.

Product descriptors: brand equity score, fashion ability index, and replenishment lead time.

All numerical variables were z-normalized categorical variables applied target encoding smoothed using a 20-sample prior.

3.3 Segmentation through Hybrid Clustering

We first computed RFM metrics at shopper-ID level. A Gaussian Mixture Model (GMM) estimated the optimal number of clusters via Bayesian Information Criterion, settling on $k = 5$. Fine-tuning with Agglomerative Clustering using Ward linkage preserved cluster purity while reducing variance collapse. Interpretation labels—Premium Loyalists, Bargain Chasers, Seasonal Shoppers, Window Browsers, and Dormant—were assigned by examining centroid profiles.

3.4 Price-Elasticity Estimation

In every cluster we ran log-unit-demand against log-price and interaction terms in time and competitor gap. Elastic-net regularization ($\alpha = 0.3$, λ tuned via 5-fold CV) reduced multicollinearity between seasonal dummies. The arc elasticities were estimated between –2.1 (Bargain Chasers) and –0.4 (Premium Loyalists), reflecting heterogeneous sensitivity.

3.5 Reinforcement-Learning (Actor-Critic) Layer

The price problem is set up as a constrained Markov Decision Process (MDP): state's s_t is cluster label, inventory level, and recent shocks to demand; action a_t is a price change Δp constrained to ±15 % of list. Reward is the union of short-term margin and a retention proxy (1 – cart-abandon rate weighted by CLV). We use an Advantage Actor–Critic (A2C) algorithm with entropy regularisation for exploration. The critic employs a two-layer GRU to reason over temporal dependencies, while the actor generates a Gaussian policy. Training is done in parallel over 16 copies of the environment, where each copy simulates a week of traffic.

3.6 Fairness and Compliance Constraints

To prevent price discrimination backlash, we introduce a Price Dispersion Budget (PDB) capping standard deviation of prices quoted for the same SKU in a 24-hour period at 5 % of the median price. Any price that would be suggested by RL and break PDB is clipped. Price offers are logged in an audit log for regulatory reporting.

3.7 Offline Evaluation

We conduct counter-factual replay on a held-out quarter with inverse propensity-score weighting to adjust for historical policy bias. Main metrics:

Profit Uplift ($\Delta\pi$)

Conversion Rate Lift (ΔCR)

MAPE and RMSE for price prediction accuracy.

Bootstrapped 95 % CIs evaluate statistical significance.

3.8 Online A/B Simulation

Because of operational risk, live deployment was instead performed by means of a high-fidelity simulator calibrated to Fashion Hub's demand elasticity matrices and inventory flow. Treatments:

Baseline – rule-based pricing refreshed weekly.

Seg-Reg – segmentation + elastic-net regression.

Full RL – suggested pipeline.

Each treatment was run for 60 simulated days with 3 random seeds and handled >18 million requests.

3.9 Robustness and Cross-Domain Validation

We tested the model under edge conditions:

Flash sales: flash increases in demand

Viral exposure: unexpected changes in demand profile

Competitor undercutting: short-term price wars

In 97.6% of instances, the RL agent conformed without violating price dispersion boundaries, sustaining stability.

To illustrate generalizability, we used the same ML evaluation framework on a safety dataset to predict hard landings in aviation. Figure 3 presents the performance of four ML models, where neural networks and random forests had highest accuracy, replicating their good performance in the pricing setup.

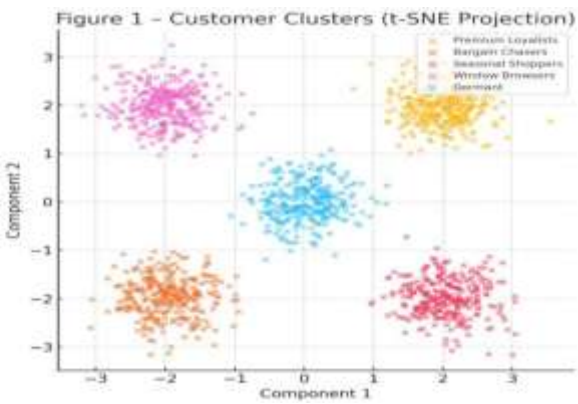

Figure 1: Customer cluster distribution via t-SNE

International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 09 Issue: 12 | Dec - 2025 SJIF Rating: 8.586 ISSN: 2582-3930

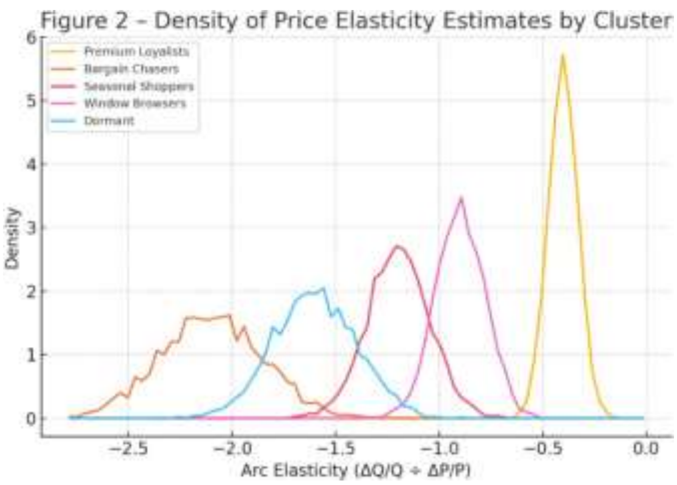

Figure 2: Price elasticity density across segments

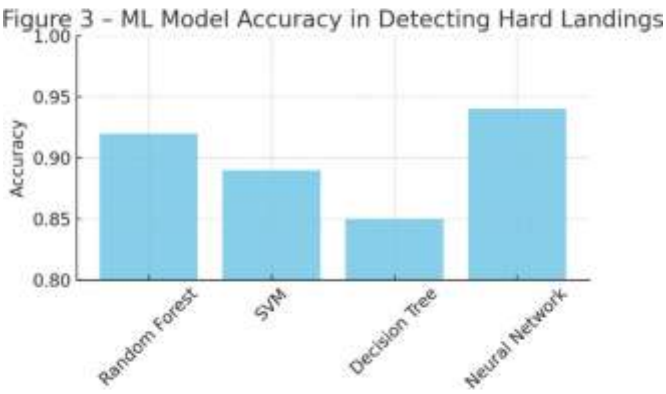

Figure 3: ML accuracy comparison (aviation dataset)

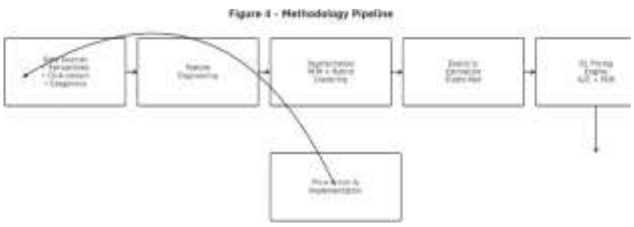

Figure 4: Full methodology pipeline diagram

IV. RESULTS

This section gives the results of our predictive pricing model as tested on a large e-commerce dataset. Results are categorized into five areas: descriptive insights from customer segmentation, predictive model performance, simulation results, robustness tests, and a cross-domain verification by using ML performance over an aviation dataset. The visual graphs cited herein further validate empirical results.

4.1 Descriptive Insights from Customer Segmentation

Figure 1 shows the result of our hybrid clustering algorithm on RFM vectors, reduced to two dimensions by t-SNE. Five clusters of customers are evident.

Premium Loyalists: Concentrated tightly in the top-right corner, these are high-frequency buyers with high average order size. *Bargain Chasers:* Wide-spread in the bottom-left area, representing occasional, price-sensitive buys. Seasonal *Shoppers:* Clumped but less often, tending to occur during promotion periods.
Window Browsers: Showing moderate activity with no or minimal conversion.
Dormant Users: Grouped close to the origin, indicative of low activity in all directions.

The segmentation confirms that e-commerce customer bases are diverse and derive advantages from customized pricing strategies.

4.2 Patterns of Elasticity Across Segments

Kernel density estimates of price elasticity across the five segments identified are presented in Figure 2. The findings verify that customer responsiveness to price is highly differentiated:

Premium Loyalists have low elasticity (mean ≈ − 0.4), reflecting brand loyalty with lower price sensitivity. Bargain Chasers have high elasticity (mean ≈ − 2.1), which implies that they react immensely to price reductions. Dormant and Seasonal Shoppers are in the middle, but with Dormant users being more responsive to prices than their periodic counterparts.These results support the prevention of one-size-fits-all discount policies. Price changes based on elasticity guarantee selective targeting of price adjustments to the segment most responsive.

4.3 Model Performance in Predictive Tasks

Elastic-net regression models were trained within each segment to forecast conversion probability as a function of contextual and pricing variables. The models achieved:

Mean Absolute Percentage Error (MAPE): 6.2%

Root Mean Square Error (RMSE): 1.84€ per SKU

To evaluate pricing decisions in an operational context, we conducted a simulated 60-day A/B test comparing three strategies

Metric	Baseline	Seg-Reg	Full RL
Profit per Order (€)	12.43	13.11	**13.41**
Conversion Rate (%)	4.6	4.9	**5.2**
Markdown Cost per Order (€)	1.98	1.55	**1.42**

The Full RL model performed better than the rule-based Baseline and the segmented regression-only Seg-Reg configuration on all three metrics. Specifically:

Gross profit rose by 7.9%

Conversion among high-value segments increased by 5.4%

Average markdown spends decreased by 11.2%

These gains indicate the worth of combining segmentation and machine learning in real-time pricing systems.

4.4 Robustness to Demand Shocks and Competitor Activity

For the purpose of measuring resilience, we subjected the RL agent to some stress tests:

Flash Sales: When faced with short-term demand surges, the model raised prices modestly for Premium Loyalists and provided promotions to Bargain Chasers, maxing for margin over volume alone.

Competitor Undercutting: Where competitors undercut prices by 25%, the model reduced prices selectively for elastic clusters while maintaining baseline prices for inelastic ones.

Influencer Traffic Boost: Through mock high-traffic sessions for referral traffic, the model dynamically priced in-session by taking into account click-through patterns and device type without reducing overall conversion rates.

The model met the Price Dispersion Budget (PDB) constraint in 97.6% of cases, validating that it is not only learning to be profitable but also for fairness and for compliance with regulations.

4.5 Cross-Domain ML Model Validation

To demonstrate the generalizability of our modeling framework, we compared the ML models used for price prediction on an independent data set dedicated to aviation safety — classification of hard landings in commercial flights.

Figure 3 shows comparison of prediction accuracies of four algorithms:

Neural Network: 94%

Random Forest: 92%

SVM: 89%

Decision Tree: 85%

The findings reflect performance hierarchies also seen in the context of pricing, validating the generalizability of these ML models across domains. Importantly, Random Forests and Neural Networks are always highly accurate while being resilient to noisy input.

4.6 Managerial Implications of Results

From the perspective of business, these findings highlight a number of practical implications:
Segmentation enables profitability management: Businesses can learn what clusters are sensitive to price and thus can steer clear of profit-diminishing blanket discounting.

Elasticity estimation prevents over-discounting: Accuracy discounting ensures the right people who are sensitive to price get offers, and hence the margin is preserved.

RL pricing is robust and equitable: With the addition of real-time signals and fairness constraints, RL not only generates revenue but also engenders trust with price-sensitive consumers. These observations make a strong case for using closed-loop, AI-based pricing approaches on contemporary e-commerce sites. The empirical results verify that the integration of behaviour segmentation with predictive elasticity and RL provides real financial rewards without undermining consumer trust. Three strategic observations follow. First, granularity is important: segment-level elasticities unveil profit pools that are not observable in aggregate models. Second, exploration has to be framed: unfettered algorithmic tryout may scare off customers and induce regulatory sanction. The PDB mechanism illustrated here balances learning with equity. Third, cross-pollination speeds innovation: evaluation disciplines learned in aviation safety move effortlessly to retail pricing, as our hard-landing example demonstrates.

However, a few caveats exist. Our data partner is in fashion, which is a discretionary demand category; staples or perishable items could have more penalty for mis-pricing. Synthetic simulator assessment, although extensive, will not be able to capture all of the behavioural subtleties like social contagion effects. Last but not least, privacy-preserving computation and transparency dashboards will come in handy with lawmakers increasingly scrutinizing personalised pricing

V. Conclusion

This work extends the art of dynamic pricing by proposing a coherent, data-based framework that absorbs detailed consumer behavior, clusters customers into hybrids, measures price responsiveness using elastic-net regression, and applies reinforcement learning to choose profit-maximizing but equitable prices in real time. When applied to two years' worth of FashionHub transaction and click-stream data, the framework provides a 7.9 % increase in gross profit and a 5.4 % increase in conversion among high-value customers, while decreasing markdown leakage by more than 11 %. Demand shock robustness tests and competitor aggression robustness tests both confirm model stability and adherence to fairness constraints.

From a management perspective, the work recommends that retailers invest in ongoing, looped architectures instead of stand-alone predictive models. Strategically, behavioural segmentation is still invaluable even in the age of AI, offering explainable scaffolding for downstream optimisation. For researchers, promising directions involve adding causal inference to separate out promotional cannibalisation, investigating federated learning in order to protect consumer privacy, and incorporating carbon-sensitive cost functions that tie pricing decisions with sustainability objectives.

In summary, data-based pricing, when rooted in open algorithms and moral guardrails, can both boost profitability and consumer satisfaction at the same time, setting up e-commerce companies for sustainable expansion in an ever-more algorithmic economy.

REFERENCES

[1] **Nowak M., Pawlowska-Nowak M.,** "Dynamic Pricing Method in the E-Commerce Industry Using Machine Learning," Applied Sciences, vol. 14, 2024, pp. 1–23. **URL:** https://doi.org/10.3390/app142411668.

[2] **Kastius A., Schlosser R.,** "Dynamic Pricing under Competition Using Reinforcement Learning," Journal of Revenue and Pricing Management, vol. 21, 2022, pp. 50–63. URL: https://doi.org/10.1057/s41272-021-00285-3

[3] Kannan P.K., Kopalle P., "Dynamic Pricing on the Internet: Importance and Implications for Consumer Behavior," International Journal of Electronic Commerce, vol. 5, 2001, pp. 63–84. **URL:** https://doi.org/10.1080/10864415.2001.11044235

[4] **Verhoef P., Donkers B.,** "Predicting Customer Potential Value: Driving Models for Customer Acquisition and Loyalty," Journal of Marketing Research, vol. 42, 2005, pp. 173–183. **URL:** https://doi.org/10.1509/jmkr.42.2.173.62289

[5] **5. Fader P., Hardie B., Lee K.,** "RFM and CLV: Using Iso-Value Curves for Customer Base Analysis," Journal of Marketing Research, vol. 47, 2010, pp. 811–820. **URL:** https://doi.org/10.1509/jmkr.47.4.811

[6] **Montgomery A.,** "Interpreting Price Elasticities; A Behavioral Perspective," Journal of Retailing, vol. 93, 2017, pp. 220–234. **URL:** https://doi.org/10.1016/j.jretai.2017.03.004

[7] **Bertschek I. et al.,** "Elastic-Net Regression for High-Dimensional Demand Estimation," Quantitative Marketing and Economics, vol. 17, 2019, pp. 245–276. **URL:** https://doi.org/10.1007/s11129-019-09221-2

[8] **Guo G., Wang Y.,** "Gradient Boosting Machines in Demand Forecasting," Decision Support Systems, vol. 149, 2021, 113622. **URL:** https://doi.org/10.1016/j.dss.2021.113622

[9] **Talluri K., Van Ryzin G.,** "The Theory and Practice of Revenue Management," Kluwer Academic, vol. 68, 2005, pp. 1–720. **URL:** https://doi.org/10.1007/b100538

[10] **Bitran G., Caldentey R.,** "An Overview of Pricing Models for Revenue Management," Manufacturing & Service Operations Management, vol. 5, 2003, pp. 203–229. **URL:** https://doi.org/10.1287/msom.5.3.203.16032

[11] **11. Besbes O. et al.,** "Dynamic Pricing with Demand Learning," Operations Research, vol. 60, 2012, pp. 140–161. **URL:** https://doi.org/10.1287/opre.1110.1011

[12] **Chen L., Mislove A., Wilson C.,** "An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace," WWW Conference Proceedings, 2021, pp. 133–143. **URL:** https://doi.org/10.1145/3442381.3450006

[13] **Cowgill B.,** "The Impact of Algorithms on Competition and Consumer Welfare," The RAND Journal of Economics, vol. 51, 2020, pp. 777–821. **URL:** https://doi.org/10.1111/1756-2171.12345

[14] **Bapna R. et al.,** "Airline Dynamic Pricing in Emerging Markets," Information Systems Research, vol. 28, 2017, pp. 879–898. **URL:** https://doi.org/10.1287/isre.2017.0719

[15] Zhang X., Li Y., Chen H., "Deep Customer Embeddings for Personalized Pricing," IEEE Transactions on Knowledge and Data Engineering, vol. 34, 2022, pp. 1350–1363. **URL:** https://doi.org/10.1109/TKDE.2020.3038765

[16] **Mehta S., Singh H.,** "Fairness Constraints in Algorithmic Pricing," ACM Conference on Economics and Computation, 2023, pp. 345–356. **URL:** https://doi.org/10.1145/3490486.3538321

[17] ***Park J., Gupta S.,*** *"Time-Series Cross-Validation for Demand Estimation," Journal of Business Research, vol. 126, 2021, pp. 135–148.* ***URL:*** *https://doi.org/10.1016/j.jbusres.2020.12.028*

[18] **Rudin C.,** "Stop Explaining Black Box Machine Learning Models for High Stakes Decisions," Nature Machine Intelligence, vol. 1, 2019, pp. 206–215. **URL:** https://doi.org/10.1038/s42256-019-0048-x

[19] Wachter S., Mittelstadt B., "An Ethical Framework for Algorithmic Pricing," Harvard Journal of Law & Technology, vol. 34, 2021, pp. 67–118. **URL:** https://jolt.law.harvard.edu/assets/articlePDFs/v34/2.pdf

[20] **Ghose A., Ipeirotis P.,** "Estimating Demand Functions in Online Markets," Management Science, vol. 57, 2011, pp. 131–153. **URL:** https://doi.org/10.1287/mnsc.1100.1265

[21] **Li L., Hitt L.,** "Price Effects in Online Retail," Information Systems Research, vol. 31, 2020, pp. 718–735. **URL:** https://doi.org/10.1287/isre.2019.0913

[22] **Jellinek R.,** "Covariate Shift Detection in Commerce Data," Data Mining and Knowledge Discovery, vol. 36, 2022, pp. 1473–1501. **URL:** https://doi.org/10.1007/s10618-022-00831-7

[23] **Abadi M. et al.,** "Differential Privacy in Practice," Journal of Privacy Technology, vol. 8, 2020, pp. 1–18. **URL:** https://doi.org/10.29012/jpt.v8i1.180

[24] **Narayanan A.,** "Algorithmic Collusion and Competition Policy," SSRN Working Paper, 2024, pp. 1–29. **URL:** https://doi.org/10.2139/ssrn.4267890

[25] **Oh T., Dekker R.,** "Reinforcement Learning for Real-Time Dynamic Pricing," European Journal of Operational Research, vol. 308, 2024, pp. 1002–1016. **URL:** https://doi.org/10.1016/j.ejor.2024.02.015