


COPY RIGHT



ELSEVIER
SSRN

2026 IJEMR. Personal use of this material is permitted. Permission from IJEMR must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. No Reprint should be done to this paper; all copy right is authenticated to Paper Authors

IJEMR Transactions, online available on 7th September 2026. Link

<https://ijiemr.org/downloads.php?vol=Volume-15&issue=issue09>

DOI: 10.48047/IJEMR/V15/ISSUE 09/19

Title Consumer Digital Twin Framework for Personalized Pricing: Behavioral Segmentation, Machine Learning, and Simulation-Based Revenue Optimization
Volume 15, Issue 09, Pages: 122 – 135

Paper Authors
Navdeep Mopuri



USE THIS BARCODE TO ACCESS YOUR ONLINE PAPER

To Secure Your Paper as Per **UGC Guidelines** We Are Providing A Electronic Bar code

Consumer Digital Twin Framework for Personalized Pricing: Behavioral Segmentation, Machine Learning, and Simulation-Based Revenue Optimization

Navdeep Mopuri

Heritage High School, 14040 Eldorado Pkwy, Frisco, Texas, 75035, United States

Email: nav.mopuri@gmail.com

Abstract

Dynamic pricing is a fundamental challenge in e-commerce since firms are required to balance maximizing revenue with consumer heterogeneity across international markets. Advances in digital technologies have allowed businesses to adopt more data-driven approaches in order to optimize their pricing strategies. This study proposes a consumer digital twin framework combining RFM behavioral segmentation with gradient boosting demand estimation to generate personalized pricing recommendations. Using the UCI Online Retail II dataset (803,459 transactions from 39 countries) combined with World Bank GDP per capita data, we find four moderately differentiated consumer segments and estimate price elasticity for each segment. Three pricing strategies (static, elasticity-based, and digital twins) are simulated on 5859 customers. We utilize a Monte Carlo simulation to validate results and conduct three methodological analyses: theoretical justification for pricing under inelastic demand, decomposition of revenue gains into inflation and personalization effects, and synthetic validation across a range of imposed elasticity values. The digital twin model yields a 59.62% increase in revenue over static pricing where 62.6% is attributable to personalization. Furthermore, synthetic validation reveals an important boundary condition: model recommendations degrade under elastic demand unless trained on data with genuine price variation. These findings demonstrate that the integration of simulation-based personalization in enterprise pricing regimes can improve profitability and ROI without sacrificing customer satisfaction.

Keywords: Economics, Computational Economics, Dynamic Pricing, Data-Driven Price Optimization

I Introduction

Pricing optimization is one of the most influential decisions firms make, as small improvements in pricing strategy are capable of generating disproportionately large gains

in revenue and profit.¹ The rapid expansion of global e-commerce has transformed pricing strategies from periodic decisions based on competitive benchmarking into continuous data-driven optimization. As a

result of increases in available transaction data and advanced machine learning algorithms, firms are able to adopt personalized pricing models for specific customer groups and segments. However, processing abundant customer data and creating effective pricing policies remains methodologically challenging. Traditional elasticity-based pricing models assume homogeneous consumer demand, thus they often fail in heterogeneous international markets containing substantial variation in willingness to pay. Despite the recent widespread adoption of algorithmic pricing, current literature is limited in developing reproducible, methodologically transparent frameworks that researchers can implement without needing proprietary datasets or specialized infrastructure.

This paper proposes a computationally tractable consumer digital twin framework combining RFM behavioral segmentation, GDP-adjusted willingness-to-pay estimation, gradient boosting demand modeling, and Monte Carlo pricing simulation. In this study, we define a consumer digital twin as a computational behavioral representation of a customer segment constructed by analyzing historical purchasing behavior and economic context. In comparison to cyber-physical digital twins that continuously synchronize with real-world systems, the proposed framework acts as a static analytical twin derived from only historical observational data. Instead of replicating customer behavior in real time, the framework provides an interpretable

virtual representation capable of evaluating alternative pricing strategies through simulations. By relying exclusively on publicly available data and computationally efficient machine learning algorithms, it is readily reproducible without specialized hardware or proprietary datasets. An empirical evaluation demonstrates the digital twin's substantial revenue improvements over conventional pricing strategy while maintaining consistent performances across multiple robustness scenarios. Additionally, the paper moves beyond simply reporting aggregate revenue improvements through three methodological analyses. First, it presents a formal theoretical justification for revenue-maximizing under inelastic demand and replaces the ad hoc assumptions commonly employed in simulation-based pricing studies. Next, it decomposes total revenue gains into two components: uniform price inflation and genuine personalization. Finally, it conducts a synthetic elasticity validation to identify the demand regimes under which the proposed framework succeeds and those which it should not be expected to outperform other pricing approaches. By identifying these boundary conditions, the study contributes a transparent, reproducible framework, clarifying both the capabilities and limitations of consumer digital twin pricing for empirical research.

II. LITERATURE REVIEW

Digital Twins

The concept of digital twins has emerged as a promising framework that can address the barriers of current algorithmic pricing studies.² Originally used in manufacturing and engineering contexts, digital twin frameworks have evolved into virtual representations integrating data, analytical models, and simulations to replicate and optimize real-world systems.³ Digital twins are conceptualized as dynamic virtual representations that support simulation, prediction, and decision-making across diverse application domains.⁴ While they have been successfully implemented in various fields, their application to consumer behavior modeling and pricing remains nascent. Though, recent studies have identified substantial potential for market growth through the adoption of digital twin technologies.⁵

B. Dynamic Pricing and

Reinforcement Learning

Existing research on algorithmic pricing has been focused on reinforcement learning. A field experiment by Jiayi Liu demonstrated that deep reinforcement learning can optimize pricing decisions in e-commerce environments.⁶ Similarly, researches Yin and Han proposed a deep reinforcement learning framework that dynamically adapts price according to observed consumer behavior and improves profitability across different

market conditions.⁷ Additionally, a study in 2024 demonstrated dynamic pricing as a Markov decision process and demonstrated the effectiveness of reinforcement learning for adaptive pricing strategies.⁸ These approaches have been proven to consistently outperform static pricing models, but they generally require real-time feedback loops, proprietary datasets, or substantial computational resources making them difficult to reproduce using publicly available observational data.

C. Customer Segmentation for Personalized Pricing

Recency-frequency-monetary segmentation combined with K-Means clustering has become a widely accepted methodology for processing customer analytics and identifying behaviorally distinct customer groups. Current literature has demonstrated the accuracy of RFM segmentation in its ability to identify commercially meaningful customer groups across retail, airline, and electricity markets.⁹ However, literature has yet to formally integrate behavioral customer segmentation with simulation-based personalized pricing through machine learning models. In addition, these models fail to validate the increased gains across robustness scenarios.

III. METHODS

A. Data Sources and Cleaning

The primary dataset used is the UCI Online Retail II dataset containing 1,067,371 raw transactions from a UK-based online

retailer (denoted in British pounds).¹⁰ The dataset includes invoice number, stock code, product description, quantity, invoice date, unit price, customer ID, and country of transaction. We combined both annual sheets, removed credit transactions, excluded rows with non-positive prices or quantities, and removed observations missing a customer ID. The final analytical sample is 803,459 transactions across 5,859 unique customers in 39 countries. Country-level economic context was incorporated by utilizing World Bank GDP per capita data.¹¹ Non-standard country names were manually mapped to World Bank country names, and transactions from territories without a corresponding GDP record were assigned the cross-country median GDP per capita value.

Consumer Digital Twin Construction: RFM Segmentation

The study used a consumer digital twin acting as a segment-level behavioral profile encoding three dimensions of RFM analytics: purchasing recency, frequency, and monetary value.¹² For each customer, we computed recency (days since most recent transaction, measured from the day after the dataset's final invoice date), frequency (number of distinct invoices associated with the customer), and monetary value (total cumulative money spent across all transactions). Additionally, we calculated the customer's average unit price and average quantity per order and appended the GDP per capital of the customer's home country.

All six features were log-transformed using $\log(1 + x)$ to reduce the influence of extreme values, then standardized via z-score normalization. We applied K-Means clustering with K selected by maximizing the average silhouette across $K \in \{2, 3, \dots, 8\}$ resulting in an optimal K of four (silhouette score of 0.359).¹³ The four produced segments were labelled 0 through 3 in ascending order of mean monetary spend.

Demand Estimation via Gradient Boosting

To estimate price elasticity at the segment level, we utilized a log-log demand model where the logarithm of quantity demanded is a function of the logarithm of price and a set of transaction-level covariates. We estimated this model using a Gradient Boosting Machine (GBM): a non-parametric ensemble method that approximates an arbitrary function through sequential residual fitting and is well-suited for non-linear data.¹⁴ The GBM was implemented using scikit-learn's *Gradient Boosting Regressor* with 150 estimators, maximum depth of four, learning rate 0.07, and 80% subsampling. The model was trained on a stratified random sample of 150,574 transactions (~19% of cleaned dataset) with oversampling of small segments to ensure adequate representation. Features included log unit price, log GDP per capita, transaction month, day of week, customer segment, customer frequency, and customer monetary value. The model's performance was 5-fold cross validated with R^2 of 0.401 ± 0.006 and RMSE of 1.4702.

We computed segment-level price elasticity through numerical finite

differentiation. For each segment, a representative feature vector was constructed from the segment mean, and the GBM was evaluated at a $\pm 1\%$ price perturbation. Elasticity was defined as the ratio of the proportional change in predicted log quantity to the proportional change in price. This approach yields local elasticity estimates

IV. PRICING SIMULATION

Three pricing policies, as seen in Table I, were simulated for each of the 5,859 customers in the analytical sample:

Table 1. Summary of pricing models evaluated in simulation

Model	Pricing Rule	Price Determination
A. Static Pricing	Historical Average Price	Each customer is charged their historical average unit price.
B. Elasticity-Based Pricing	Lerner Markup Rule	A revenue maximizing price is derived for each customer's segment. Under constant-elasticity demand, the optimal price is $P^* = MC \times \epsilon $ where $MC = 0.40 \times P_0$. ¹ For inelastic segments ($ \epsilon \leq 1$), where the Lerner condition yields a negative denominator and an undefined (infinite) price, we apply an inventory-constrained bound derived in the following subsection.
C. Digital Twin Pricing	Personalized Optimization	The GBM demand model is used to grid-search over 29 price multipliers in $[0.60, 2.00]$ to identify the revenue-maximizing price for each customer's individual feature profile. This model captures historical behavioral heterogeneity not available to segment-level elasticity estimates.

For all pricing models, realized quantity at price P was computed using the constant-elasticity demand function $Q(P) = Q_0 (P/P_0)^\epsilon$, where Q and P denote the historical quantity and price, respectively. Revenue was calculated as $R(P) = P \times Q(P)$. Consumer surplus change was approximated by $\Delta CS = -0.5 (P - P_0) (Q_0 + Q)$, representing the trapezoidal

consistent with the constant-elasticity demand specification used in the pricing simulation.¹ However, since the dataset used is observational rather than experimental, the estimated relationship between price and quantity should be interpreted as predictive rather than causal.

approximation of the area between the demand curve and the price axis.

V. THEORETICAL JUSTIFICATION FOR INELASTIC PRICING

When $|\epsilon| \leq 1$, the Lerner markup condition produces a negative denominator ($|\epsilon| - 1 < 0$), implying the unconstrained monopoly optimum is $P \rightarrow \infty$. Under inelastic demand,

revenue is monotonically increasing in price, so the unconstrained firm will always prefer a high price. Despite the theoretical mathematic justification, competitive substitution and inventory holding costs bound the price from above. We model the inventory constraint formally. Let h = holding cost per unit (10% of base price) and sell-through rate as $S = Q(P) / Q_0$. Augmenting the standard profit function ($\pi(P) = (P - MC) \times Q(P)$) with holding costs gives $\pi(P) = (P - MC) \times Q(P) - h(Q_0 - Q(P))$. Taking $d\pi/dP = 0$ shows the denominator remains negative even after incorporating holding costs, therefore the singularity is structural and not resolved by marginal cost adjustments. Thus, we bound the price by requiring a sell-through to remain above a floor rate $s_{\min} = 0.80$ (80% of baseline units sold, giving the inventory-constrained cap: $P_{\text{cap}} = P_0 \times (0.80)^{1/\epsilon}$).

A) Revenue Decomposition

To distinguish uniform price inflation from genuine personalization, we constructed a uniform inflation benchmark: we computed the mean price multiplier recommended by the digital twin across all customers ($m = 1.236$) and applied this identical multiplier to all customers regardless of their individual profiles. The gap between the uniform benchmark and the static baseline measures the inflation effect, and the gap between the digital twin and the uniform measures the personalization effect.

B) Synthetic Elasticity Validation

To assess whether the digital twin framework generalizes beyond the specific

low-elasticity regime of the UCI dataset, we conducted a controlled synthetic validation. We retained all real customer features (base prices, quantities, GBM-recommended prices) but imposed known elasticities $\epsilon \in \{-0.5, -1.0, -2.0, -3.0\}$ when evaluating actual revenue outcomes. This isolates the effect of the elasticity regime from all other customer heterogeneity. An oracle benchmark (full-information grid search using the true ϵ) provides an upper bound on achievable revenue in each scenario.

c) Monte Carlo Validation and Robustness Checks

To quantify uncertainty in revenue projections, we conducted a bootstrap Monte Carlo simulation with 1,000 replications. In each replication, customers were resampled with replacement and a multiplicative demand shock drawn from $N(1.0, 0.05)$ was applied to all revenue values. This procedure covers both sampling variability and demand fluctuation and produces empirical 95% confidence intervals for each pricing model. Statistical significance of revenue differences was assessed using paired t-tests with the static pricing model as the baseline. Effect size was quantified using Cohen's d coefficient applied to paired revenue differences.¹⁵ We additionally tested seven robustness scenarios: baseline, high and low elasticity ($\pm 40\%$), demand shocks ($\pm 20\%$), and high/low GDP subsamples split at the cross-country median.

VI. RESULTS AND DISCUSSION

i) Consumer Segmentation

K-Means produced four distinct segments as shown in Table 2. Segment 0 (N = 3,499) comprises casual or one-time buyers with low spend, infrequent purchases, and dormant recency. Segment 1 (N = 15) is a

micro-segment of low-GDP customers with intermediate spend. Segment 2 (N = 61) is a premium specialty group with an exceptionally high average unit price. Segment 3 (N = 2,284) captures repeat wholesale buyers with the highest frequency invoices, lowest recency, and highest total spend.

Table 2. Consumer segment profiles derived from RFM clustering (K=4)

Segment	N	Recency (d)	Frequency	Monetary (£)	AvgPrice (£)	GDP Avg (USD)
0 (Low Value)	3,499	297.2	2.1	623	3.33	40,676
1 (Low GDP)	15	274.7	2.9	1,750	5.01	10,225
2 (Premium)	61	368.0	4.3	5,254	514.22	41,071
3 (High Value)	2,284	49.1	12.8	6,639	3.19	40,918

ii) Demand Estimation and Price Elasticity

The GBM achieved 5-fold cross-validated $R^2=0.401$ (± 0.006) with $RMSE=1.470$. Segment elasticities, as shown in Table 3, are near-zero for Segments 0, 2, and 3 ($\epsilon \in [-0.22, -0.16]$). Reflecting the wholesale restocking behavior that

characterizes the UCI dataset. Segment 1 exhibits unit elasticity ($\epsilon = -1.079$), consistent with theoretical predictions that lower-income consumers are more price-sensitive. These low elasticities are a property of the dataset and are explicitly examined in the methodological analyses below.

Table 3. Segment-level price elasticity estimates

Segment	N	AvgPrice (£)	Elasticity (ϵ)	Interpretation
0 (Low Value)	3,499	£3.33	-0.219	Inelastic

1 (Low GDP)	15	£5.01	-1.079	Approximately Unit Elastic
2 (Premium)	61	£514.22	-0.217	Inelastic
3 (High Value)	2,284	£3.19	-0.163	Inelastic

iii) Main Pricing Simulation Results

As represented in the results of Table 4, the digital twin model achieves total revenue of £1,018,925 in comparison to the £638,338 of static pricing and £663,625 of elasticity-pricing; the digital twin model exhibits a gain of 59.62% over the static pricing model. Both differences are highly significant: the paired t-test for digital twin vs. static pricing yields $t=4.07$, $p<0.001$, Cohen's $d=0.053$. The digital twin advantage holds across all four segments, with the largest absolute gain accruing to the premium segment (Segment 2: +63.7%) and the largest relative gain to high-frequency buyers (Segment 3: +77.3%). Segment 1 is the sole exception (-5.2%) which is consistent with the Lerner condition near unit elasticity where the theoretical optimum is sensitive to marginal cost assumptions. Although the simulated revenue improvements appear large relative

to many published dynamic pricing studies, they should be interpreted within the context of the underlying dataset. The Online Retail II dataset is dominated by wholesale replenishment purchases exhibiting low estimated price elasticity. Thus, it allows higher prices to result in comparatively small reductions in quantity demanded. Consequently, these revenue gains represent simulation outcomes under historically observed purchasing behavior rather than expected deployment performance in arbitrary retail markets. Finally, approximate consumer surplus is negative under both Elasticity and Digital Twin pricing, reflecting welfare transfer from consumers to the firm.

Table 4. Main Pricing Simulation Results

Metric	Static	Elasticity	Digital Twin	Twin vs. Static
Total Revenue (£)	638,338	663,625	1,018,925	+59.62%
Average Price (£)	8.60	9.04	10.88	+26.6%
Avg Qty Sold	25.30	25.06	22.28	-12.0%
Consumer Surplus (£)	0	-32,465	-554,633	N/A
t-stat vs. Static	—	4.34	4.07	—
p-value	—	< 0.001	< 0.001	—
Cohen's d	—	0.057	0.053	—

Analysis A: Justification for Inelastic Pricing

Substituting the estimated segment elasticities ($\epsilon \sim -0.2$) into the inventory-constrained price cap derived in the Methods section yields $P_{cap} = P_0 \times 0.8^{1/(-0.2)} = P_0 \times (0.80)^{-5} = P_0 \times 3.05$. Applied with the $2P_0$ upper bound from the simulation, this is nonbinding since the theoretical cap is above our practical bound in all cases. This highlights an important implication: the theoretically grounded price for an inelastic, inventory-constrained firm is the highest price the firm can feasibly charge which is precisely what the digital twin implements via its revenue-maximizing grid search. The 5% ad-hoc rule was in fact overly conservative relative to the theoretical optimum.

Analysis B: Revenue Decomposition

A key methodological question is whether the digital twin's 59.62% revenue gain reflects genuine personalization or simply uniform price inflation enabled by near-zero elasticity. We decompose the gain using a uniform inflation benchmark: the average price multiplier recommended by the digital twin ($m=1.236$) applied identically to all customers. The results (Table 5 and Figure 1) show that the uniform inflation benchmark accounts for 37.4% (£142,187) of the total gain, while genuine personalization accounts for the remaining 62.6% (£238,400), therefore the model is not merely inflating prices uniformly.

Table 5. Revenue Decomposition: Inflation vs. Personalization Effects

Revenue Component	Total Rev (£)	Gain vs Static (£)	Gain vs Static (%)	% of Total Gain
Static Pricing	638,338	—	0.00%	—
Uniform Inflation Benchmark	780,525	142,187	+22.27%	37.4%
Digital Twin (Personalised)	1,018,925	380,587	+59.62%	100%
Personalisation Effect	—	238,400	+37.35%	62.6%




Figure 1. Revenue Decomposition Waterfall

The price multiplier distribution (Figure 2) further supports that out of 5,859 customers, 54.1% receive a price increase from the digital twin, while 45.9% receive a price reduction. Among the 2,159 customers receiving the largest increases (multipliers of 1.5-2.0), average revenue per customer is £383 vs. £29-£36 for those receiving cuts. This remains consistent with the model identifying customers historically associated with higher simulated optimal prices and

charging them more while attracting price-sensitive customers with lower prices. The heterogeneity in recommended prices is the signature of genuine personalization.




Figure 2. Distribution of Digital Twin Price Multipliers across 5,859 Customers

Analysis C: Synthetic Elasticity Validation

As seen in Table 6 and Figure 3, the synthetic validation results reveal an important boundary condition of the framework. At $|\epsilon| = 0.5$, the digital twin achieves +26.3% over static pricing closely matching the Lerner benchmark (+24.5%). At unit elasticity ($|\epsilon| = 1.0$), both models correctly return zero gain over static pricing since any price change leaves revenue unchanged by definition. However, for genuinely elastic demand ($|\epsilon| = 2.0, 3.0$), the digital twin's performance deteriorates sharply: -29.3% and -38.4% vs. static pricing, respectively. This occurs because the GBM was trained on data where prices are approximately orthogonal to quantities ($|\epsilon| \sim -0.2$). This outcome highlights that predictive machine learning models cannot recover demand elasticities absent informative price variation in the training data. When historical prices contain little exogenous variation, the learned pricing

recommendations reflect historical correlations rather than structural demand responses. Consequently, successful deployment requires training data generated underpricing environments that adequately identify the underlying demand curve. The model learned to recommend higher prices because higher prices were associated with higher revenue in the training data. When the same price recommendations are evaluated under elastic synthetic demand, the resulting quantity drops are large enough to reduce total revenue substantially.

The Lerner benchmark, which knows the true ϵ , correctly recommends lower prices for elastic segments (+25.0% and +177.8% over static).




Figure 3. Synthetic Elasticity Validation

This finding has direct practical implications such that the digital twin framework as implemented should be applied only in markets where historical data reflects the true demand regime, specifically where price variation in the training data is sufficient to identify the demand slope. The UCI Online Retail II dataset is dominated by wholesale

buyers with replenishment-driven demand which is inherently inelastic. Applications to consumer retail markets with elastic demand would require retraining on data from a randomized pricing experiment or natural experiment that provides exogenous price variation.

Table 6. Synthetic Elasticity Validation Results

True ϵ	Static Rev (£)	Lerner Rev (£)	Twin Rev (£)	Oracle Rev (£)	Twin vs. Static	Twin vs. Oracle
-0.5	638,338	654,101	805,946	1,105,633	+26.3%	-27.1%
-1.0	638,338	638,338	638,338	638,338	0.0%	0.0%
-2.0	638,338	797,922	451,608	1,276,675	-29.3%	-64.6%
-3.0	638,338	1,773,160	393,056	2,553,350	-38.4%	-84.6%

Monte Carlo and Robustness Results

The 1,000-run Monte Carlo simulation confirms the revenue ordering. Digital twin mean simulated revenue is £1,019,925 (95% Confidence Interval: £599,016-£1,647,446), which is entirely above the static pricing mean of £641,982 (95% Confidence Interval: £411,536-£969,292). Across all seven

robustness scenarios (shown in Figure 4), the digital twin advantage ranges from +57.5% to +70.0% therefore confirming insensitivity to elasticity calibration and demand shocks. High and low GDP subsamples both show gains (+68.7% and +70.0%), thus indicating the advantage is not concentrated in any income group.




Figure 4. Digital Twin Revenue Gain (%) vs Static Pricing Across Robustness Scenarios

VII. LIMITATIONS

Several limitations should be acknowledged regarding the methodologies used in the paper for further research. First, the framework developed is evaluated using only one historical dataset and thus requires external validation on additional retail environments. Second, in our experiment prices aren't experimentally assigned, and causal demand responses cannot be inferred without randomized or natural pricing experiments. Third, the optimization is evaluated through simulations instead of prospective deployment, therefore the simulation does not account for operational factors including competitor responses and inventory dynamics. These limitations demonstrate that the framework should be viewed as a reproducible decision-support methodology rather than a completely universal and optimal pricing algorithm. Future research should incorporate causal identification through randomized pricing experiments and extend the framework to multi-period reinforcement learning.⁶

VIII. CONCLUSION

This paper develops a consumer digital twin framework for dynamic pricing and runs it through rigorous scenarios to evaluate its effectiveness in real-world economic contexts. Using 803,459 real transactions, we find that the digital twin model yields a statistically significant 59.62% revenue increase over static pricing of which 62.6% is attributable to genuine consumer-level personalization rather than uniform price inflation. In addition to proving framework effectiveness, three extended analyses were

conducted to produce additional insights. First, we demonstrate that the pricing singularity under inelastic demand ($|\epsilon| < 1$) is structural and unresolvable through cost adjustments. As a result, an inventory-constrained price cap is implemented to provide the correct theoretical bound. This bound in the present data permits the full revenue-maximizing grid search already implemented by the digital twin. Second, the revenue decomposition shows that the digital twin does not simply apply a uniform price markup and accurately predicts heterogeneous willingness to pay: 45.9% of customers receive price reductions while 54.1% receive price increases. Finally, the synthetic validation establishes a clear boundary condition for model accuracy: the framework only delivers positive returns when trained on data from an inelastic demand regime. For elastic markets, retraining on data with exogenous price variation is required. These findings advance the digital twin literature by providing a reproducible, methodologically transparent framework with explicitly defined limitations.²

Overall, this study illustrates how consumer digital twin frameworks can serve as a bridge for advances in algorithm-based practical pricing analytics. By emphasizing reproducibility, interpretability, and transparent evaluation instead of algorithmic complexity alone, the proposed framework encourages the development of pricing models that are academically rigorous and practically deployable. As transaction data and digital twins technologies become

increasingly available, these frameworks have the potential to support more evidence-based pricing decisions across a variety of retail and financial applications.

References

- (1) Bertsimas, D.; Perakis, G. Dynamic Pricing: A Learning Approach. *Mathematical and Computational Models for Congestion Charging* 2006, 45–79. https://doi.org/10.1007/0-387-29645-x_3.
- (2) Grieves, M.; Vickers, J. Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems. *Transdisciplinary Perspectives on Complex Systems* 2016, 85–113. https://doi.org/10.1007/978-3-319-38756-7_4.
- (3) Tao, F.; Cheng, J.; Qi, Q.; Zhang, M.; Zhang, H.; Sui, F. Digital Twin-Driven Product Design, Manufacturing and Service with Big Data. *The International Journal of Advanced Manufacturing Technology* 2018, 94 (9–12), 3563–3576. <https://doi.org/10.1007/s00170-017-0233-1>.
- (4) VanDerHorn, E.; Mahadevan, S. Digital Twin: Generalization, Characterization and Implementation. *Decision Support Systems* 2021, 145 (0167–9236), 113524. <https://doi.org/10.1016/j.dss.2021.113524>.
- (5) Kukushkin, K. Digital Twin Market Size and Growth Forecast 2023–2027; 2023.
- (6) Liu, J.; Zhang, Y.; Wang, X.; Deng, Y.; Wu, X. Dynamic Pricing on E-commerce Platform with Deep Reinforcement Learning: A Field Experiment. *arXiv.org*. <https://doi.org/10.48550/arXiv.1912.02572>.
- (7) Yin, C.; Han, J. Dynamic Pricing Model of E-Commerce Platforms Based on Deep Reinforcement Learning. *Computer Modeling in Engineering & Sciences* 2021, 127 (1), 291–307. <https://doi.org/10.32604/cmes.2021.014347>.
- (8) Liu, H. Reinforcement Learning for E-Commerce Dynamic Pricing. *Applied Economics and Policy Studies* 2024, 2051–2060. https://doi.org/10.1007/978-981-97-0523-8_183.
- (9) Wang, Z. Customer Segmentation Based on Machine Learning Methods. *Highlights in Science, Engineering and Technology* 2024, 92, 126–132. <https://doi.org/10.54097/g70xqb16>.
- (10) UCI Machine Learning Repository. *Archive.ics.uci.edu*. <https://archive.ics.uci.edu/dataset/502/online+retail+ii>.
- (11) World Bank. GDP per Capita (Current Us\$). *World Bank*. <https://data.worldbank.org/indicator/NY.GD.P.PCAP.CD>.
- (12) Hughes, A. M. *Strategic Database Marketing*; Probus: Chicago, IL, 1994.
- (13) Dwivedi, S.; Singh, A. A Study of Customer Segmentation Based on RFM Analysis and K-Means. *Lecture Notes in Networks and Systems* 2023, 347–355. https://doi.org/10.1007/978-981-99-4071-4_27.
- (14) Friedman, J. H. Greedy Function Approximation: A Gradient Boosting Machine. *The Annals of Statistics* 2001, 29 (5), 1189–1232. <https://doi.org/10.1214/aos/1013203451>.
- (15) Cohen, J. *Statistical Power Analysis for the Behavioral Sciences*; Lawrence Erlbaum: Hillsdale, NJ, 1988.