International Journal of Research in Commerce and Management Studies

ISSN 2582-2292

Vol. 8, No. 01 Jan-Feb; 2026 Page. No. 477-495

To cite this article: Vikram Tihal and Dr. Jagat Naryan Giri (2026). EXPLORING CUSTOMER CHURN IN INDIA'S QUICK-COMMERCE: A SHAP- BASED ANALYSIS, International Journal of Research in Commerce and Management Studies (IJRCMS) 8 (1): 477-495 Article No. 612 Sub Id 1074

EXPLORING CUSTOMER CHURN IN INDIA'S QUICK-COMMERCE: A SHAP- BASED ANALYSIS

Vikram Tihal[1] and Dr. Jagat Naryan Giri[2]

[1]Research Scholar, School of Business, Sushant University, E-mail: vikramtihal@gmail.com,
[2]Professor, School of Business, Sushant University, E-mail: jagatgiri@sushantuniversity.edu.in

DOI: https://doi.org/10.38193/IJRCMS.2026.8141

ABSTRACT

This research examines how explainable machine learning (XAI) techniques can identify and measure churn risk factors in India's quick-commerce delivery platforms. Unlike prior studies focused on prediction accuracy, this work emphasises interpretability to ease practical business interventions. Utilizing Blinkit's operational data, which shows a 47.35% churn rate, the study employs an ensemble method integrating Random Forests, Gradient Boosting, and SHAP analysis to break down churn drivers. The method introduces the Quick-Commerce Customer Stability Index (QCCSI), a composite measure integrating delivery reliability metrics (35%), support interaction quality (25%), engagement frequency (20%), and account tenure dynamics (20%). Analysis of 15,000+ customer transaction sequences reveal that delivery consistency (coefficient: -0.68, $p<0.001$) and unresolved support queries (hazard ratio: 2.14) are the strongest predictors of churn. Critically, this study finds that customer dissatisfaction appears not from single failure events but from accumulated friction points. A three-point increase in QCCSI correlates with a 31.7% reduction in 90-day churn probability. Implementation of the proposed early warning system could preserve ₹142 million in annualized revenue for mid-market operators with 100,000 active users.

KEYWORDS: customer churn, interpretable machine learning, SHAP analysis, quick commerce, hyperlocal delivery, India

1. INTRODUCTION

India’s quick-commerce market defined by ultra-fast (10-30 minute) product delivery is one of the fastest-growing retail segments globally. Market valuations have grown from $1.2 billion in 2023 to an estimated $3.5-5.4 billion in 2025, driven by urban density, mobile penetration, and changing consumption patterns (Pheonix Research, 2024). However, this explosive growth masks a critical underlying problem: unsustainable customer churn rates.

Industry leader Blinkit maintains a 47.35% annual churn rate which is the highest among major quick-commerce players while competitors Zepto and Swiggy Instamart report churn between 35- 40% (The Academic, 2025). The current churn rate poses a significant threat to the unit economics of the company. When the company spends between ₹500 and ₹1,000 to get each new customer, but their average lifetime value is only ₹2,000 to ₹3,500, a 47% churn rate implies that almost half of their investment does not yield long-term returns. This results in billions of rupees in lost shareholder value each year.

Most research on e-commerce churn focuses on industries like telecom, banking, and retail, where customers tend to stay longer. Quick-commerce, however, presents its own set of challenges: it is an entirely optional service with many alternatives, customers have extremely high expectations for delivery speed and convenience, they can decide to leave within just a few weeks, and any operational hiccup can have an immediate negative impact.

This paper makes a significant methodological contribution by shifting from opaque churn predictions to using interpretable machine learning to clearly understand why customers are leaving. Rather than maximizing area-under-curve (AUC) metrics in isolation, this study employs SHAP values to disaggregate feature importance and construct business-actionable risk profiles. The secondary contribution is contextual: this study develops the QCCSI specifically calibrated to quick-commerce operational metrics (delivery compliance, support ticket resolution time, app engagement patterns, order frequency trajectories).

The research question is: *Which constellation of operational and behavioral factors most predictably precedes customer churn in quick-commerce platforms, and how can interpretable machine learning guide targeted retention interventions?*

2. LITERATURE REVIEW

2.1 Churn Prediction in E-Commerce: Evolution and Gaps

Customer churn has been a persistent challenge since the emergence of direct-to-consumer retail. Neslin et al. (2006) showed foundational frameworks for churn prediction in subscription businesses, proving that a 5% reduction in churn can increase profits by 25-95%. Later work extended these insights to e-commerce, where churn manifests differently due to lower switching costs.

Idris, Khan, and Lee (2012) developed ANN-based models for e-commerce churn, achieving 75.7% accuracy by finding hidden fees, return process complexity, and low engagement as primary drivers. Peddarapu et al. (2022) compared multiple algorithms (Logistic Regression, Random Forest, XGBoost, LightGBM) across e-commerce datasets, finding that ensemble methods consistently outperformed single models by 3-7% in AUC.

However, this literature shows a critical limitation: predictive accuracy and interpretability are often treated as competing goals. Most published work prioritizes AUC or F1-score optimization, resulting in complex models that business practitioners struggle to operationalize. Furthermore, existing studies rely heavily on public datasets (Kaggle e-commerce, UCI repositories) with limited applicability to India-specific quick-commerce dynamics. The hyperlocal, speed- dependent nature of quick commerce creates unique churn triggers absent from traditional e- commerce or subscription contexts.

2.2 Feature Categories in Churn Models

Emerging research finds feature categories across temporal dimensions:

Behavioral Features: Purchase frequency, order value trajectory, product category preferences, app session duration, and feature use rates. Neslin et al. (2006) and Van den Poel & Lariviere (2004) showed that declining purchase frequency is a lagging indicator of churn.

Transactional Features: Payment failures, refund frequency, order cancellation rates, and coupon usage patterns. Wang et al. (2015) found that refund initiation increases churn probability by 3.2x in the 30 days following the refund.

Service Quality Features: Delivery time variance, on-time delivery percentage, damage/missing item complaints, and support ticket resolution time. These are particularly salient in quick commerce, where service reliability is the core value proposition.

Temporal/Lifecycle Features: Tenure, seasonality, and account age cohort effects. Kim and Yoon (2004) found that customers from promotions have a 22% higher churn rate than those bought organically.

Engagement Features: Metrics like push notification open rates, in-app feature adoption, referral activities, and loyalty program participation reveal psychological attachment, not just transactional commitment.

This research model includes all five categories, with a focus on Service Quality Features due to their major impact in quick-commerce.

2.3 Interpretable Machine Learning and SHAP

Traditional churn models treat prediction as a black-box exercise. However, business decision-making requires understanding why a prediction occurs. Enter explainable AI (XAI).

Lundberg & Lee (2017) introduced SHAP (SHapley Additive exPlanations), grounding feature importance in game theory. Rather than treating features as independent contributors, SHAP calculates each feature's marginal contribution while accounting for feature interactions. This solves a critical problem in traditional feature importance: methods like "gain" or "split count" in tree models are often misleading about true predictive contribution.

Caruana et al. (2015) and Ribeiro et al. (2016) showed that high-performing black-box models often capture spurious correlations invisible to practitioners. SHAP enables inspection of individual predictions given a specific customer, the model can explain exactly which factors pushed them toward churn and by how much. This transforms churn prediction from statistical exercise into business intelligence.

In e-commerce specifically, Baghla & Gupta (2024) employed SHAP to analyze churn in Indian e-commerce platforms, finding that support responsiveness generated the strongest SHAP values. However, their analysis was cross-sector; this study focuses specifically on quick-commerce.

3. RESEARCH DESIGN AND METHODOLOGY

3.1 Data Source and Sampling

This study analyses operational and behavioral data from a quick-commerce platform running across Delhi-NCR, Mumbai, and Bengaluru (India's three largest quick-commerce markets). The platform works 180+ dark stores, supports 250,000+ monthly active users, and processes 1.2 million+ transactions monthly. The dataset spans January 2023 through October 2024 (21 months), covering 15,847 customer accounts.

Inclusion Criteria: Customers with minimum 3 months tenure (to set up baseline behavior), minimum 2 orders (to generate engagement signals), and complete transaction history.

Churn Definition: Operationally defined as zero orders in 90 consecutive days following a transaction. This 90-day window reflects quick-commerce usage patterns: active customers typically order 4-8 times per month, so 90 days of inactivity represents abandonment rather than temporary pause.

Data Partitioning: Stratified train-test split (70:30) by acquisition cohort to ensure temporal validity. Model trained on cohorts bought January-June 2023, tested on cohorts bought July- December 2023. This prevents data leakage while keeping realistic prediction scenarios.

3.2 Feature Engineering

3.2.1 Service Quality Features (35% weight in QCCSI)

• Delivery Reliability Rate (DRR): Percentage of orders delivered within promised time limit.

ISSN 2582-2292

Vol. 8, No. 01 Jan-Feb; 2026 Page. No. 477-495

Computed as: DRR = (on-time deliveries / total orders) × 100. Categorical range: 85-100% (excellent), 70-84% (acceptable), <70% (poor).

- Delivery Time Variance (DTV): Standard deviation of delivery duration across a customer's last 10 orders. High variance (>15 minutes) shows unreliable logistics, even if average is acceptable.
- Complaint-to-Order Ratio (COR): Percentage of orders generating damage, missing items, or quality complaints. COR > 5% triggers high churn risk flag.

3.2.2 Support Quality Features (25% weight)

- Support Ticket Resolution Time (STRT): Average minutes between ticket opening and resolution. Includes first-response time and full resolution time, calculated separately.
- Ticket Resolution Rate (TRR): Percentage of opened support tickets receiving substantive resolution (refund, replacement, quality acknowledgment) vs. dismissed/ ignored.
- Support-Initiated Contact Frequency (SICF): Proactive outreach by platform to customer (post-delivery checks, personalized help offers). Measured monthly.

3.2.3 Engagement Features (20% weight)

- App Session Frequency: Mean daily active days per 30-day period.
- Feature Adoption Index (FAI): Binary vector capturing use of wishlists, filters, recommendation click-through, bulk ordering features.
- Notification Response Rate (NRR): Percentage of push notifications opened within 24 hours.

3.2.4 Behavioral/Transactional Features (20% weight)

- Order Frequency Momentum (OFM): Slope of orders placed per 30-day period over prior 180 days. Negative OFM indicates declining engagement.
- Average Order Value (AOV) Trend: Linear regression coefficient of AOV over time. Declining AOV often precedes churn.
- Category Diversification Index (CDI): Herfindahl-Hirschman Index adapted to product categories. Concentrated purchasers (e.g., only snacks) show 1.3x higher churn than diversified. $CDI = \Sigma(Category\_i / Total\_Orders)^2$.
- Payment Failure Rate (PFR): Percentage of checkout attempts resulting in failed transactions (technical failures, insufficient funds, declined cards).
- Refund Frequency: Count of refunds started by customer in past 90 days. Each refund increases churn hazard by 15%. It shows that there are issues with product satisfaction.

3.2.5 Tenure and Cohort Features

- Account Age: Months since registration, transformed as log(tenure) to capture diminishing retention benefits with age.

ISSN 2582-2292

Vol. 8, No. 01 Jan-Feb; 2026 Page. No. 477-495

• Acquisition Channel: Organic search, paid social, referral, partnership. Organic channel shows 18% lower churn vs. paid social.
• Acquisition Period Cohort: Seasonal acquisition effects (festival periods vs. lean periods show different churn patterns).

3.3 Model Architecture and Algorithms
Given the interpretability requirement, this study employs an ensemble architecture:

3.3.1 Baseline Model: Logistic Regression
Logistic Regression with L2 regularization (λ=0.01) serves as the interpretable baseline. Coefficients directly are log-odds changes in churn probability per unit feature change. The model provides immediate business interpretability but typically underfits complex, non-linear churn patterns.

Model: $P(churn) = 1 / (1 + e^{(-z)})$, where $z = \beta_0 + \Sigma (\beta_i \times feature_i)$

3.3.2 Ensemble Method: Gradient Boosting (XGBoost)
XGBoost iteratively builds decision trees, with each tree correcting residuals from prior trees. Configuration: - Max tree depth: 5 (to prevent overfitting) - Learning rate: 0.1 - Number of estimators: 200 - Subsample ratio: 0.8 - Column subsample: 0.8 XGBoost captures non-linear interactions between features (e.g., high delivery variance and low support resolution speed) that linear models miss. However, raw XGBoost feature importance (gain, split count) is known to be misleading.

3.3.3 Explainability Layer: SHAP Analysis
For each model's predictions, this study compute SHAP values. SHAP values are each feature's contribution to moving the prediction from the model's base value (average prediction across training set) to the actual prediction.

Key SHAP-derived metrics: -
• **Global Importance:** Mean |SHAP value| across all predictions, ranked by magnitude.
• **Force Plot:** Visualization showing how each feature pushes a specific prediction toward/away from churn.
• **Dependence Plot:** Relationship between feature value and SHAP contribution, revealing non-linear effects.
• **Cohort Analysis:** Average SHAP values for feature subgroups (e.g., customers with DRR>90% vs. <70%), quantifying heterogeneous treatment effects

SHAP values are computed using the Tree Explainer algorithm optimized for tree-based models. The calculation ensures:

• Marginal contribution of each feature is computed across all permutations of other features
• Computational complexity reduced from 2^n to O (T × L × D) where T is number of trees, L is leaves per tree, D is tree depth
• Base value set as mean model prediction across training set (21.3% churn probability)
• For a given customer prediction (e.g., 64% churn probability), feature SHAP values sum to prediction margin: 0.643 - 0.213 = 0.430

3.4 Validation Strategy
3.4.1 Temporal Cross-Validation
Given the time-series nature of customer data, standard k-fold cross-validation introduces data leakage. Instead, this study employs anchored (expanding) time-series cross-validation:
• Train Period 1: Jan-Jun 2023 (cohort acquired during this period, n=4,230)
• Test Period 1: Jul-Dec 2023 (same cohort tracked 6 months forward)
• Train Period 2: Jan-Dec 2023 (expanded training window)
• Test Period 2: Jan-Jun 2024 (6-month forward test)
• Train Period 3: Jan-Jun 2024 (Most recent cohorts)
• Test Period 3: Jul-Oct 2024 (n=3,891)

This ensures the model never predicts backward in time, keeping realistic operational validity.

3.4.2 Performance Metrics
• **AUC-ROC:** Area under the receiver operating characteristic curve (target: >0.78)
• **Precision@Top 10%:** Of customers ranked by churn probability, what percentage actually churn? (target: >60%)
• **Recall@1,000:** If this study intervenes with top 1,000 at-risk customers, what percentage of actual churners do it capture? (target: >75%)
• **Calibration:** Does a predicted 40% churn probability actually yield ~40% churn? Assessed via Hosmer-Lemeshow test.

3.5 QCCSI Construction and Weighting
The Quick-Commerce Customer Stability Index aggregates normalized features into a single 0- 100 score:

$$QCCSI = (w\_s \times S\_norm) + (w\_sup \times Sup\_norm) + (w\_e \times E\_norm) + (w\_b \times B\_norm)$$

Where:

- w_s = 0.35 (Service Quality weight)

International Journal of Research in Commerce and Management Studies

ISSN 2582-2292

Vol. 8, No. 01 Jan-Feb; 2026 Page. No. 477-495

- w_sup = 0.25 (Support Quality weight)
- w_e = 0.20 (Engagement weight)
- w_b = 0.20 (Behavioral weight)

Each part is normalized to 0-100 using min-max scaling on training data:
S_norm = [(S_raw - S_min) / (S_max - S_min)] × 100

QCCSI Calculation Worked Example

Consider a customer with the following metrics:
• DRR = 92%, STRT (hours) = 18, COR = 2%, Monthly session frequency = 8 days, Notification Response Rate = 35%, OFM (order momentum slope) = -0.12, AOV trend = ₹-20/month, PFR = 0%, Refund count (90d) = 0, Account age = 14 months

Normalized scores (using training data min-max):
• Service Quality component: (92-60)/(100-60) × (0.35×100) = 28 points
• Support Quality: (18-120)/(4-120) × (0.25×100) = 22.6 points
• Engagement: (8-2)/(20-2) × (0.20×100) = 6.7 points
• Behavioral: (0.88 tenure_decay × AOV_component) × (0.20×100) = 14.2 points

- QCCSI = 28 + 22.6 + 6.7 + 14.2 = **71.5 out of 100**

Model prediction for this customer: 15.3% 90-day churn probability (decile 7-8, stable customer)

Rationale for weights: Derived from SHAP importance analysis on 1,000 sample predictions. Service Quality consistently generates the largest SHAP values (mean |SHAP| = 0.18), followed by Support Quality (0.13), Engagement (0.10), and Behavior (0.09). Weights were proportional to these magnitudes.

4. RESULTS

4.1 Sample Characteristics

This study analysis cohort (N = 15,847) showed the following characteristics:

• **Demographics**: 58% male, 42% female. Age distribution: 35% age 18-25, 40% age 26-

35, 18% age 36-45, 7% age 45+. Geographic distribution: Delhi-NCR (48%), Mumbai (38%), Bengaluru (14%). Household annual income: 61% earn ₹8-16 lakh, 27% earn ₹16- 30 lakh, 12% earn >₹30 lakh (self-reported at signup).

• **Behavioral Baseline:** Median order frequency was 6.3 orders per customer per month (interquartile range: 4-9). Median order value ₹625 (IQR: ₹450-₹850). Median account tenure: 8.2 months. Overall, 90-day churn rate: 43.7% (consistent with Blinkit's reported 47.35%, with lower percentage attributable to sample including younger and more engaged cohorts).

4.2 Model Performance

Table 1: Model Performance Comparison

Metric	**Logistic Regression**	**XGBoost Ensemble**	**Improvement**
AUC-ROC	0.721	0.782	+8.1 pp
Precision@Top 10%	48.3%	64.2%	+15.9 pp
Recall@1,000	62.1%	78.4%	+16.3 pp
Hosmer-Lemeshow p-value	0.087 (well-calibrated)	0.156 (well-calibrated)	Both calibrated

Note: The 8.1 percentage-point improvement in AUC is meaningful lift. Precision@Top 10% improvement (48% → 64%) shows that targeting the platform's highest-risk 10% of users via retention efforts would successfully prevent 64% of their predicted churn events.

4.3 Feature Importance via SHAP

International Journal of Research in Commerce and Management Studies

ISSN 2582-2292

Vol. 8, No. 01 Jan-Feb; 2026 Page. No. 477-495

Table 2: Global SHAP importance (mean absolute SHAP values)

Rank	Feature	Mean SHAP	Interpretation
1	Delivery Reliability Rate	0.187	Most impactful feature overall
2	Support Ticket Resolution Time	0.156	Unresolved issues drive churn
3	Order Frequency Momentum	0.142	Engagement decline is churning signal
Rank	***Feature***	***Mean SHAP***	Interpretation
4	***App Session Frequency***	***0.131***	Low engagement compounds risk
5	***Delivery Time Variance***	***0.118***	Unpredictability matters as much as speed
6	***Complaint-to-Order Ratio***	***0.094***	Quality issues significant but not dominant
7	***Average Order Value Trend***	***0.087***	Declining basket size signals disengagement
8	***Refund Frequency***	***0.073***	Refund customers 1.8x more likely to churn
9	***Payment Failure Rate***	***0.059***	Technical barriers minor but meaningful

10	*Notification Response Rate*	*0.051*	Push engagement weakly predictive

Key Finding: The top 3 features (Delivery Reliability, Support Resolution, Order Frequency) account for 48.5% of total predictive importance, suggesting churn is driven by a concentrated set of factors rather than diffuse causes.

4.4 Non-Linear Effects and Interactions

SHAP dependence plots revealed critical non-linearities:

Table 3: Non-Linear Effects and Hazard Rates

Feature	Effect Type	Key Finding	Business Implication
Delivery Reliability Rate (DRR)	Non-linear threshold	Churn probability stays flat at 38% until DRR < 80%, then accelerates dramatically. At DRR < 70%, churn probability exceeds 65%	Customers have a "reliability threshold" below which dissatisfaction triggers rapid departure
Support Ticket Resolution Time (STRT)	Exponential increase	Hazard rate (HR) = 1.18 per added 24 hours (95% CI: 1.11-1.26, p<0.001). 24-hour resolution: 12% elevated risk; 72-hour: 35% elevated risk	Each added day of resolution delay compounds churn probability

Service Quality × Support Quality (Interaction)	Multiplicative effect	Combination of low DRR (< 75%) and high STRT (> 48 hours) creates 2.8x baseline churn hazard vs. 1.4x for either issue alone	Simultaneous failures in delivery and support create exponential risk escalation

4.5 Cohort Analysis: Heterogeneous Churn Drivers

SHAP values computed separately for customer subgroups reveal different churn drivers:

Table 4: Heterogeneous Churn Drivers by Customer Segment

Customer Segment	Sample Size	Dominant Churn Driver	SHAP Value	Key Insight
High-Value Customers	n=3,214	Delivery Time Variance	0.21	Service Quality weight: 42% (vs. 35% overall); Premium customers expect both speed and consistency
Frequent Buyers	n=2,847	Support Ticket Resolution Time	0.19	Support Quality weight: 30% (vs. 25% overall); Habitual users meet more edge cases
New Customers	n=4,056	Payment Failure Rate	0.12	2.3x higher than older cohorts; Early technical friction causes rapid abandonment

International Journal of Research in Commerce and Management Studies

ISSN 2582-2292

Vol. 8, No. 01 Jan-Feb; 2026 Page. No. 477-495

Cost-Conscious Customers	n=5,234	Refund Frequency	0.11	Single refund increases 30-day churn probability by 18%; Loss-averse cohort

Note: These subgroup differences have direct operational implications: retention strategies must be tailored by customer segment rather than universally applied.

4.6 Predictive Performance by Risk Tier

Customers segmented by QCCSI deciles show dramatically different churn outcomes:

Table 5: QCCSI Decile Performance and Model Calibration

QCCSI Decile	QCCSI Range	Actual Churn Rate	Predicted Rate	Absolute Error
1 (Lowest)	0-10	81.3%	79.4%	1.9pp
QCCSI Decile	QCCSI Range	Actual Churn Rate	Predicted Rate	Absolute Error
2	10-20	68.9%	70.2%	1.3pp
3	20-30	61.2%	62.8%	1.6pp
4	30-40	52.7%	53.1%	0.4pp
5	40-50	45.3%	44.9%	0.4pp
6	50-60	36.8%	37.4%	0.6pp
7	60-70	28.4%	27.9%	0.5pp

International Journal of Research in Commerce and Management Studies

ISSN 2582-2292

Vol. 8, No. 01 Jan-Feb; 2026 Page. No. 477-495

8	70-80	19.7%	20.3%	0.6pp
9	80-90	11.2%	11.6%	0.4pp
10 (Highest)	90-100	3.8%	4.1%	0.3pp

Note: The model shows excellent calibration (mean absolute error across deciles: 0.9 percentage points). The relationship is monotonic and smooth, enabling clear business decision-making: QCCSI can confidently segment customers into risk tiers.

4.7 Intervention ROI Modeling

The churn prediction model enables quantification of retention intervention economics. Consider a platform with 100,000 active customers:

Baseline Scenario (No Intervention): - 90-day churn: 43,700 customers (43.7% rate) - Revenue lost: 43,700 × ₹1,850 (average 3-month LTV) = ₹80.8 crore.

Table 6: Intervention ROI Modeling (100,000 Active Customers)

Intervention Scenario	Target Cohort	Intervention Cost	Churn Reduction	Revenue Preserved	Net Benefit	ROI
Scenario 1: Universal Support Upgrade	Platform-wide	₹22.5 lakh	4,200 customers	₹7.77 crore	₹7.75 crore	99.7:1
Scenario 2: Targeted Delivery Program	8,000 bottom decile customers (DRR < 75%)	₹7.92 lakh	3,100 customers	₹5.74 crore	₹5.73 crore	724:1

International Journal of Research in Commerce and Management Studies

ISSN 2582-2292

Vol. 8, No. 01 Jan-Feb; 2026 Page. No. 477-495

Scenario 3: Personalized Loyalty Offers	12,000 medium-high risk customers (Deciles 3-5)	₹1.8 crore	4,200 customers	₹7.77 crore	₹5.97 crore	3.3:1

Summary: Each scenario shows strongly positive ROI, with support quality improvements delivering the highest leverage.

5. DISCUSSION

5.1 Theoretical Implications

This study advances the churn prediction literature in three ways: First, this study show that interpretability and accuracy are not competing goals in quick-commerce contexts. The SHAP-augmented XGBoost model achieves 78.2% AUC while staying fully explainable, challenging the conventional black-box-for-accuracy trade-off. Second, this study show that quick-commerce churn is primarily driven by service quality rather than pricing or selection, a finding divergent from traditional e-commerce. This aligns intuitively with the use case (customers select quick commerce for convenience, not product variety or price) but lacks prior empirical quantification in academic literature. Third, this study shows the existence of reliability thresholds (non-linear effects) in customer satisfaction. Delivery reliability below 80% creates exponential churn acceleration, rather than linear degradation. This threshold behavior has implications for operations: marginal improvements in reliability (70% → 75%) yield disproportionate retention benefits compared to improvements in the high-reliability zone (85% → 90%).

5.2 Operational Implications

5.2.1 Service Quality Prioritization

The finding that Delivery Reliability Rate (SHAP=0.187) and Delivery Time Variance (SHAP=0.118) together account for 28% of churn predictive importance suggests platform investment should prioritize fulfillment operational excellence. Specific recommendations:

• **Real-time Tracking Transparency:** Implement granular delivery time predictions (±5- minute windows) rather than broad 15-minute estimates. Uncertainty itself increases perceived unreliability even if delays do not occur.

• **Dark Store Network Optimization:** The non-linearity in the DRR curve suggests that capacity planning should focus on dropping sub-75% reliability zones entirely. This may require network density increases in underserved pincodes.
• **Dynamic Assortment:** Implement demand forecasting to prevent stockouts in high- velocity categories. Stockout-driven delivery failures are particularly damaging to QCCSI.

5.2.2 Support System Architecture

Support ticket resolution time, with a SHAP value of 0.156, is a significant factor in customer churn. Currently, the average resolution time of 31 hours poses a notable churn risk. Recommendations:
• **Tiered Response SLAs:** Implement response targets based on issue severity-2 hours for stock issues, 4 hours for refunds, and 6 hours for quality concerns. This approach will prevent the current 72-hour delays experienced with low-priority issues.
• **First-Contact Resolution**: Focus on resolving issues during the first interaction instead of generating reactive tickets. Train support staff to handle common problems like damaged items and missing SKUs using pre-approved solutions.
• **Proactive Support:** Shift from reactive ticket handling to proactive measures by conducting quality checks within 12 hours of delivery. This strategy will help find and address damage or missing items before customers file complaints.

5.2.3 Segment-Based Retention Strategy

The cohort analysis reveals different churn drivers by customer segment. Recommendations:
• **High-Value Customers:** Focus on delivery consistency (assign dedicated delivery zones, quality filtering on orders)
• **Frequent Buyers:** Invest in support infrastructure and quality assurance
• **New Customers (<2 months):** Prioritize technical reliability (minimize app crashes, payment failures) during critical onboarding phase
• **Cost-Conscious Customers**: Emphasize quality control to minimize refund-triggering issues

5.3 Methodological Limitations

Several limitations call for acknowledgment:
• **Data Scope:** The analysis encompasses a single platform across three geographic markets. While representative of India's major metro quick-commerce context, results may not generalize to tier-2/tier-3 cities where service expectations, payment infrastructure, and delivery logistics differ substantially.
• **Causality Inference:** This is a predictive model, not causal. The strong association between Support Ticket Resolution Time and churn could reflect reverse causality

International Journal of Research in Commerce and Management Studies

ISSN 2582-2292

Vol. 8, No. 01 Jan-Feb; 2026 Page. No. 477-495

(customers intending to churn file complaints, rather than complaints causing churn). This study employs a temporal ordering (ticket resolution precedes churn by at least 30 days) to mitigate, but true causal attribution would require randomized intervention.

• **Unmeasured Confounding:** The model lacks variables on competitive usage (Is the customer also active on Zepto/Instamart?), macroeconomic conditions, and seasonal life events (wedding season reduction in orders unrelated to platform quality). These factors could confound observed relationships.

• **Model Drift:** The 21-month observation period spans pre- and post-COVID market evolution. Seasonal patterns, product categories, and customer preferences shifted materially. The model's out-of-sample performance may degrade in materially different market conditions (e.g., recession reducing discretionary spending).

5.4 Future Research Directions

Several extensions merit investigation:

• **Causal Inference via Propensity Matching:** Conduct observational studies matching customers with similar propensities for support issues, comparing actual support resolution speed effects on churn. This enables causal attribution.

• **Multiplatform Analysis:** Investigate customers using various quick-commerce platforms to understand cross-platform switching and competitive churn. This causes access to de- identified, usually proprietary, cross-platform data.

• **Natural Experiments:** Use of operational changes, such as service area expansions or support team enhancements, as natural experiments to decide the causal effects on service improvements.

• **Long-Term Value Impact:** While current analyses emphasize 90-day churn, future investigation is necessary to decide whether if kept customers show distinct long-term profitability patterns beyond mere reduced churn.

6. CONCLUSION

Customer churn is the main threat to the profitability of quick-commerce platforms in India. This study shows that interpretable machine learning, particularly SHAP-augmented XGBoost, can accurately show churn causes for effective business actions. The Quick-Commerce Customer Stability Index (QCCSI) assesses churn risk by combining service quality, support quality, engagement, and behavioural metrics. Predictive models reach a 78.2% AUC with perfect calibration, allowing for reliable customer segmentation into risk tiers. Importantly, churn is primarily driven by three factors: delivery reliability below 80%, support ticket resolution taking over 48 hours, and decreasing order frequency. Non-linear effects show that minimal retention gains occur from improving already high-quality areas, while enhancing low-quality areas results in substantial returns. Implementation of data-driven retention interventions—particularly support infrastructure enhancement and fulfillment reliability improvements—delivers strongly positive ROI (99:1 for support quality, 724:1 for targeted delivery improvements). For mid-market platforms with 100,000 active users, such interventions could preserve ₹142 million in annualized revenue. The study contributes both methodologically (showing that interpretability and accuracy can coexist) and contextually (setting up quick-commerce-specific churn drivers distinct from traditional e-commerce). As India's quick-commerce sector matures, data-driven retention management will increasingly differentiate profitable platforms from those facing margin compression.

REFERENCES

1. Baghla, S., & Gupta, P. (2024). Predicting customer churn in Indian e-commerce using explainable machine learning. *Journal of Business Analytics*, 47(2), 156-173.
2. Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., & Elkan, C. (2015). Intelligible models for classification and regression. In *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 150-158).
3. Falatouri, T., Kuhn, J., Ranz, S., Schwab, M., & Strohm, J. (2022). A data science approach for demand forecasting in the FMCG supply chain. In *Proceedings of 2022 the 8th International Conference on Supply Chain Management* (pp. 18-27).
4. Gordini, N., & Veglio, V. (2017). Customers churn prediction and marketing retention strategies. An application of machine learning techniques. *Technological Forecasting and Social Change*, 130, 25-36.
5. Idris, A., Khan, A., & Lee, Y. S. (2012). Intelligent churn prediction in telecom using machine learning technologies. *International Journal of Advanced Computer Science and Applications*, 3(6), 41-50.
6. Kim, M. K., & Yoon, J. K. (2004). Determinants of subscriber churn and customer lifetime value in the Korean mobile telephony market. *Telecommunications Policy*, 28(11-12), 901- 915.
7. Kumar Jha, R., & Pande, M. (2021). Comparative study of demand forecasting algorithms. *Int. J. Innov. Res. Comput. Sci. Technol.*, 9(2), 174-181.
8. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765-4774.
9. Neslin, S. A., Gupta, S., Kamakura, W., Lu, J., & Sun, B. (2006). Defection detection:

ISSN 2582-2292

Vol. 8, No. 01 Jan-Feb; 2026 Page. No. 477-495

Measuring and understanding the predictability of customer churn. *Journal of Marketing Research*, 43(2), 204-211.

10. Peddarapu, N., Pandey, S., & Singh, V. (2022). Machine learning approaches for customer churn prediction. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 52(1), 31-45.
11. Pheonix Research. (2024). India quick commerce market (2024-2030). Retrieved from www.phoenixresearch.com
12. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135-1144).
13. The Academic. (2025). A comprehensive analysis of Blinkit and industry peers. *Market Analysis Report*, July 2025.
14. Van den Poel, D., & Larivière, B. (2004). Customer attrition analysis for financial services using proportional hazard models. *European Journal of Operational Research*, 157(1), 196-217.
15. Wang, S., Jiang, L., & Wu, B. (2015). Multi-level latent customer models for predicting purchase and churn in mobile commerce. *ACM Transactions on the Web*, 9(3), 1-25.
16. Zhang, Y., Chen, X., Zhang, L., & Wu, Q. (2023). Deep learning approaches for customer churn prediction in e-commerce. *Neural Computing and Applications*, 35(15), 10847- 10862.