International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(4s) 2026 pp. 992-1006
https://doi.org/10.70917/ijcisim-2026-2612
Article

Data Driven approach for investigation of Food delivery order cancellations on Zomato platform and development of predictive model to determine order cancellation beforehand.

Bhisaji Surve[1], Jayesh Manjarekar[2], Harshall Gandhi[3]

[1]Prin L N Welingkar Institute of Management Development & Research, PGDM, Mumbai,
Email id - aba9767@gmail.com
[2]Prin L N Welingkar Institute of Management Development & Research, PGDM, Mumbai
Email id -jjmanjrekar@gmail.com
[3]Prin L N Welingkar Institute of Management Development & Research, PGDM, Mumbai
Email id -harshil.gandhi@welingkar.org

Abstract: The reasons behind the cancellation of food orders on the Zomato platform in India are examined, based on a dataset of 12,079 food orders placed in 76 cities between January 2023 and October 2024. The study concludes that there is an overall cancellation rate of 32.3%, resulting in revenue loss and inefficiency of operations. Although promotional codes, free delivery offers and price range seem to have an impact, it is not significant in the context of the cancellations that took place. The research is descriptive analysis of several factors that affect the cancellation of orders and appraises the importance of each factor. It also suggests a machine learning based predictive model and a multi-stage analysis process to deal with cancellation issues. The results offer some practical suggestions for platform operators, such as criteria for restaurant engagement, operator capacity building and real-time assessment of cancellation risk when placing an order.
Keywords: food delivery, order cancellation, Zomato, predictive modelling

1. Introduction

The Indian market for online food delivery services has seen significant growth, and Zomato has risen to become one of its leading providers. Considering increasing digital ordering among cities and semi-cities in India, the efficiency and dependability of the delivery process have become highly important. Order cancellation has been noted as one of the most impactful service quality problems where the loss happens for the restaurant, loss of efficiency for the agent, and dissatisfied customers.

There has been considerable literature related to the issue of food delivery cancellation in India focusing mainly on demand prediction and optimization of delivery time (Kumar et al., 2021; Sharma & Patel, 2022). However, research into the issue of structural factors that influence the process of order cancellation has yet to be conducted.

This paper addresses the following core research problem:

What structural and operational factors drive food delivery order cancellations on the Zomato platform, and can these factors be used to predict cancellations at order placement time?

1.1 Research Objectives

- Identify the statistically significant predictors of food delivery cancellation across order, restaurant, and agent dimensions.
- Calculate the value of restaurant type, city tier and delivery agent performance on cancellation variance.
- Translate findings into actionable platform policy recommendations.

1.2 Scope and Limitations

The dataset includes 76 Indian cities and orders placed on Zomato from January 2023 to October 2024. This is an observational, cross-sectional analysis and causal inference is limited due to a lack of experimental variation. There is no order level field to record cancellation reasons which means that it is not possible to distinguish supply side from demand side cancellations. The restaurant type missingness (2.8%) could cause some level of bias in type-level comparisons.

2. Literature Review

The literature, nonetheless, covered within this study can be divided into four domains, which are closely related to each other. 1) the evolution and dynamics of the global food delivery market and in India.

2) the phenomenon of order cancellations in on-demand services

3) the performance and heterogenization of delivery agents in last mile logistics

4) the use of machine learning methods such as gradient boosting and explainable AI in operational risk prediction.

The research hypotheses put forward in this paper are informed by each body of literature and the analytical propositions are informed by each body of literature as well.

2.1 The emergence of online food delivery services.

The online food delivery is one of the fastest growing in the digital economy. According to the Taiwan Ministry of Digital Affairs (2024), the global food delivery market is estimated to exceed USD 449 billion by 2025, and will continue to expand at a CAGR of 10% to 20%.The Taiwan Ministry of Digital Affairs (2024) projects the food delivery market in the world is expected to reach USD 449 billion by 2025, rising at a rate of 10% to 20% CAGR. The Indian market, in particular, has experienced a significant growth, to be valued at USD 45.15 billion in 2024, and the market is projected to grow by CAGR of 22.18% up to 2034, fuelled by the adoption of smart phones, low-cost mobile data and higher disposable income, particularly in Tier-2 and Tier-3 cities (IMARC Group, 2025).

In India, the market is dominated by two platform to consumer food delivery aggregators, Swiggy and Zomato, who have a combined market share of more than 90% of the organized food delivery market (Nexdigm, 2025). According to Datum Intelligence (2025), Zomato has become profitable in FY2024 and currently has around 57% order value market share compared to Swiggy's 43%. Platform to consumer (P2C) models, which involve an aggregator taking care of the entire logistics journey, have emerged as the top position, making up approximately 65% of the market (Nexdigm, 2025). This is an aggregator model that establishes a three-sided market between consumers, restaurants and delivery agents, where the interdependencies are more complicated than in previous studies on bilateral markets.

Platform activity is unevenly distributed geographically. The high population density of restaurants, high adoption of digital ordering and established logistics in major metros like Bengaluru, Delhi-NCR, Mumbai, Hyderabad, and Chennai propel them from the top of the list by volume (Research and Markets, 2023). Yet the strategic front line has moved to smaller cities: Order volumes in Tier-2 and Tier-3 markets have doubled in a few years and cloud kitchen formats are proving to be more cost-effective (around 800 daily orders, compared to 1,300 in metros) to expand there (IMARC Group, 2025). This geographic variation is part of the data to be understood when seeking to understand variation in cancellation rates across our data.

2.2 Order Cancellation in On-Demand Delivery Services

Order cancellation is one of the least studied and most important phenomena in the field of on-demand delivery studies. Consumers in these situations are time-sensitive and impatient, and as emphasized by the paper mentioned that cancellation is a common and serious reaction to perceived service failures (Bai et al., 2024, ScienceDirect). Most importantly, the existing models have been somewhat simplistic, and have either considered that cancelation is a random event, or that it is triggered by some threshold on the wait time, or that it is modelled by an exogenous probability distribution, none of which models the structural and operational drivers found in practice (Bai et al., 2024).

In their study on the Chinese food delivery app Meituan and Swiggy, He, Liu and Liang (2024) used logistic regression, propensity score matching (PSM) and back-propagation neural networks to examine cancellation predictors. While their key insight that longer waiting times are associated with higher cancellation rates is consistent with our data with the dimension of delivery time, it's not enough: We also know that there is a significant variation in cancellation rates by type of restaurant, which is a dimension that isn't covered by their study. This indicates that there is an additional dimension (supply side characteristics: operational readiness of restaurants) to the analysis that prior research has not separated.

Deliverability time forecasting is a related problem, in which Salari, Liu and Shen (2022) pointed out that rather than the expectation of deliverability time, uncertainty of it can lead to adverse consumer outcomes, such as cancellation. The tree-based distributional forecasting proposed by them is directly applicable to their deployment architecture proposed in Section 6: delivery time uncertainty estimates can be added to a real-time cancellation risk score to give rise to composite risk signals. Liu et al. (2022) also found that there are meaningful differences in cancellation behaviour between the order matching stage (consumer-platform) and the pick-up stage (platform-restaurant), which we are unable to observe directly but led us to the hypothesis that restaurant type (proxy of restaurant-staged operational maturity) is a significant predictor.

Within the wider e-commerce literature, the issues of order cancel and price sensitivity, product substitutability, and consumer commitment signals have been explored. Lower-friction payment methods (cash and easily-reversed digital payments) are consistently linked with higher cancellation rates (Chen & Xu, 2018). Support for this directionality is weak in our data but still present with the cancelling of cash orders at 34.1% and net banking at 30.4%. It does not seem to be a strong signal, however — cash seems to be a secondary driver.

2.3 The Indian Food Delivery Context: City-Tier Infrastructure Gaps

One key aspect of cancellation rates in the city tier of our initial analysis is based on the structural asymmetry we found in the Indian food delivery literature. Key factors that give Metro markets an advantage include the density of agents, well-established restaurant onboarding protocols and shorter median delivery distances. Tier-2 and Tier-3 cities, on the other hand, have seen significant growth in their order volumes, but not seen the same proportional increase in delivery infrastructure, making agents scarce, and restaurants less ready to deliver (Research and Markets, 2023; IMARC Group, 2025).

This infrastructure deficit poses a direct challenge for cancellation. If there are few agents, they can also be located in remote areas, and orders may be placed by such agents, which can result in longer pick-up times and either a restaurant's or a platform's decision to cancel. If the standards for restaurants that are onboarding are lower (as is the case in markets that are growing rapidly), then a restaurant may be able to take delivery orders, even though they are not able to service them in the kitchen. In December 2024, WAAYU, India's first Zero Commission Food Delivery App, which was launched in Hyderabad with plans to expand into Tier-2 markets, explicitly identified restaurant onboarding quality as a key operational challenge in the smaller markets (WAAYU, 2025). This strengthens our claim that secondary cities, which have recently been expanding by Zomato, like Gorakhpur, Salem and Shimla, have cancellation rates 10-15 percentage points higher than that of metros.

Payment infrastructure in the Indian context also has an impact on the consumer behaviour aspect. As of late 2024, India is one of the most digitally connected consumer markets in the world with 886 million internet users and 99.5% of youth adopting UPI payments (Nexdigm, 2025), but consumer commitment suggested by the payment method they choose is still what it is. As expected, following the commitment-friction hypothesis from the e-commerce literature in general, the results indicate that the highest cancellation rates in our data set are for cash orders, which have no commitment to make and are immediately reversible at no cost.

2.4 Machine Learning for Operational Risk Prediction in Food Delivery

Since the year of 2020, the adoption of machine learning in the field of operational risk classification in the food delivery sector has been significantly increasing. In this area, gradient boosting methods such as XGBoost and LightGBM are the key tools that have been widely used to tackle tabular classification problems and consistently outperformed logistic regression and random forests in structured operational data (Li et al., 2021; Zhang et al., 2022).

XGBoost has several benefits for this application: it is able to process mixed feature types (categorical, continuous, binary); it naturally supports regularisation to avoid overfitting when dealing with small samples; and it is scalable and efficient for real-time deployment.

The interpretability aspect of machine learning models that are used in operational contexts has been gaining attention. Lundberg and Lee (2017) presented SHapley Additive exPlanations (SHAP), a single framework based on cooperative game theory that gives both global feature importance rankings and local (per-prediction) explanations. SHAP values are now the norm for operational deployments of tree-based models (Lundberg & Lee, 2017) and have been proven in various domains such as demand forecasting, fraud detection, and delivery time prediction. In the context of food delivery cancellations, SHAP offers the tool to convert a predictive model into an operational rule: If SHAP value for 'restaurant type = microbrewery' is high, there is a clear, human-interpretable reason to investigate that order for proactive cancellation.

He, Liu and Liang (2024) used logistic regression and neural networks for predicting food delivery cancelations and concluded that logistic regression had competitive accuracy and gave interpretable estimates of the coefficients. Yet, the linear decision boundary of logistic regression is not likely to capture the interaction effects between the restaurants, city tier and the performance of the agents that our exploratory analysis indicates are present. Logistic regression, on the other hand, cannot model non-linear interactions and threshold effects (such as 'cancellation risk rises when restaurant type is nightlife AND city tier is Tier-2 or below').

Based on the meta-analytic review by Shankar et al. (2024) of 85 studies concerning online food delivery consumer behaviour, the delivery reliability factor, defined as the degree of consistency in which the orders are completed without cancellation, is the one key factor that increases the likelihood of platform continuance intention. This places the issue of cancellation reduction on the agenda not only from an operational efficiency perspective, but also from a business perspective in relation to customer lifetime value and platform retention; this reinforces the business case for the predictive modelling framework presented in this paper.

2.5 Research Gaps and Positioning

Despite the growing literature across each of these four domains, several critical gaps remain unaddressed. First, no study has examined the specific relationship between restaurant type — particularly nightlife-oriented categories such as microbreweries, clubs, and pubs — and food delivery cancellation rates. These venues are prevalent in Zomato's restaurant catalogue but have been excluded from prior analyses that focus on quick service and casual dining categories. Second, city-tier heterogeneity in cancellation has not been empirically quantified at the granular level enabled by our 76-city dataset: prior Indian market analyses have described infrastructure gaps qualitatively without linking them to measurable cancellation outcomes. Third, agent-level cancellation analysis has been conducted in the ride-hailing context (e.g., Castillo et al., 2022) but not specifically in the food delivery context with Indian gig workers, where different incentive structures and employment conditions apply. Fourth, the prediction has not been applied to the food delivery cancellation classification problem at the point of order placement in the Indian market.

This paper directly addresses each of these gaps: it quantifies the restaurant type effect with category-level cancellation rates across 12+ venue types; it maps city-tier cancellation variation across 76 cities; it profiles agent heterogeneity across 899 agents; and it proposes and validates predictive model framework for real-time cancellation risk scoring. Together, these contributions advance the empirical literature on food delivery operations and provide immediately actionable intelligence for platform management.

3. Data Description

Outcome Variable: Order Status

The outcome variable is a three-category field: Delivered (33.9%), In-Transit (33.8%), and Cancelled (32.3%). For modelling point of view, the cancellation prediction task is outlined as a binary classification: Cancelled vs. Not Cancelled (Delivered + In-Transit).

Order Status	Count	Percentage	Implication
Delivered + In-Transit	4,093+4,089=8182	67.7%	Successful

(non-cancelled)			
Cancelled	3,897	**32.3%**	Lost revenue

Table 1: Order status distribution (n = 12,079)

The 32.3% cancellation rate is distinctly above industry standards reported in similar markets (typically 10–20%), making this a critical operational problem need of dedicated analysis.

3.1 Key Predictor Variables

The given data set having 33 columns can be grouped into three categories of predictors:

- **Order-level:** Order Amount (INR), Delivery Time, Number of Items, Delivery Distance, Payment Method, Promo Code Applied, Has Free Delivery, Tip Amount, Order Date.
- **Restaurant-level:** Restaurant Type, City, State, Cuisines, Average Cost for Two, Price Range, Aggregate Rating, Votes, Photo Count, Delivery availability.
- **Agent-level:** Delivery Agent ID (from which agent-specific cancellation rate, average delivery time, and average tip can be derived).

4. Research Methodology

4.1 Research Framework

The methodology proceeds in five subsequent stages:

- Feature Engineering
- Statistical significance testing
- Exploratory data analysis
- Agent clustering
- Predictive modelling
- Policy translation.

Each stage generates outputs that serve into the next.

4.2 Feature Engineering

Several features are drawn from raw columns to enhance predictive utility:

- **Temporal features:** Hour of day, day of week, and month extracted from Order Date. Weekend indicator (Friday–Sunday) constructed to capture nightlife effect.
- **Agent features:** Per-agent cancellation rate, mean delivery time, and mean tip amount computed from historical orders within the dataset (leave-one-out to avoid target leakage).
- **Cuisine flags:** Cuisines field parsed into binary indicators for the 15 most frequent cuisine types (North Indian, Chinese, Continental, etc.).
- **City tier:** Cities manually classified into Metro (Mumbai, Delhi, Chennai, Bengaluru, Kolkata, Hyderabad), Tier-1, Tier-2, and Tier-3 based on population and Zomato market maturity.

4.3 Statistical significance Testing

Prior to modelling, the statistical significance of candidate predictors is assessed:

4.3.1 **Chi-square tests** of independence: applied to categorical predictors (restaurant type, city tier, payment method, promo code, free delivery) against the binary cancellation outcome.

Variable	Chi-Square (χ^2)	p-value	DOF	Result
Restaurant Type	53.99	0.00043	24	Significant
City	91.46	0.095	75	Not Significant
Payment Method	4.55	0.337	4	Not Significant
Promo Code Applied	0.16	0.688	1	Not Significant
Free Delivery	0.24	0.627	1	Not Significant

Restaurant Type vs Cancellation; $p < 0.05 \rightarrow$ Significant

Table 2: Chi-square tests results

Conclusion: The relationship between type of restaurant and cancellation of orders is statistically significant. Preparation time may vary from one restaurant to another, such as casual dining restaurants vs quick service restaurants because of manpower shortage or material ready. Both are more likely to experience delays resulting in higher cancelation rates. This is one of the important findings.

Interpretation:

- No location dependence of cancellation behaviour
- Proposes uniform standards of service in cities
- The cancellation is not dependent on the payment method.
- Promotional offers do not impact on cancellation rates.
- Deliveries do not appreciably affect cancellations.

Actionable insight for Zomato platform point of view:

- **Focus on:**
  - Restaurant onboarding quality
  - Evaluate Kitchen efficiency
  - Preparation time optimization
- **Less focus needed on:**
  - Payment method policies
  - Promotional strategies for reducing cancellations

4.3.2. One-way ANOVA: For continuous predictors (order amount, delivery distance, delivery time, number of items) use One-way ANOVA (Cancelled vs. Not Cancelled groups. To determine if there are any differences between the means of continuous variables across:

Cancelled Orders (1), Non-Cancelled Orders (0)

Variable	Mean (Not Cancelled)	Mean (Cancelled)	F-Statistic	p-value	Significant (Bonferroni)
Delivery Time (mins)	47.8	54.6	18.72	0.00002	Yes

Delivery Distance (km)	4.98	5.02	0.21	0.646	No
Order Amount (INR)	812	798	1.45	0.228	No
Number of Items	3.21	3.28	0.89	0.345	No

Table 3: ANOVA Results Table

Significance Level

- **Base: α = 0.05; Number of tests = 4**
- **Bonferroni corrected α = 0.05 / 4 = 0.0125**

Only p-values < **0.0125** are considered significant.

Conclusion: The delivery time for cancelled orders is statistically significantly different from delivery time for non-cancelled orders p = 0.00002 < 0.0125 So, the average delivery time of cancelled orders is higher.

The results highlight the importance of efficient deliveries in last-mile logistics. Order characteristics (value and size) do not seem to have a major effect on the cancellation behaviour, but the delays of the delivery time are sufficient to significantly raise the delivery order failure rate. This underscores the value of logistics optimization in eCommerce operations, particularly for the time-sensitive nature of the eCommerce market.

4.4 Exploratory data analysis

Exploratory Data Analysis (EDA) was conducted to understand the structure, distribution, and relationships within the dataset.

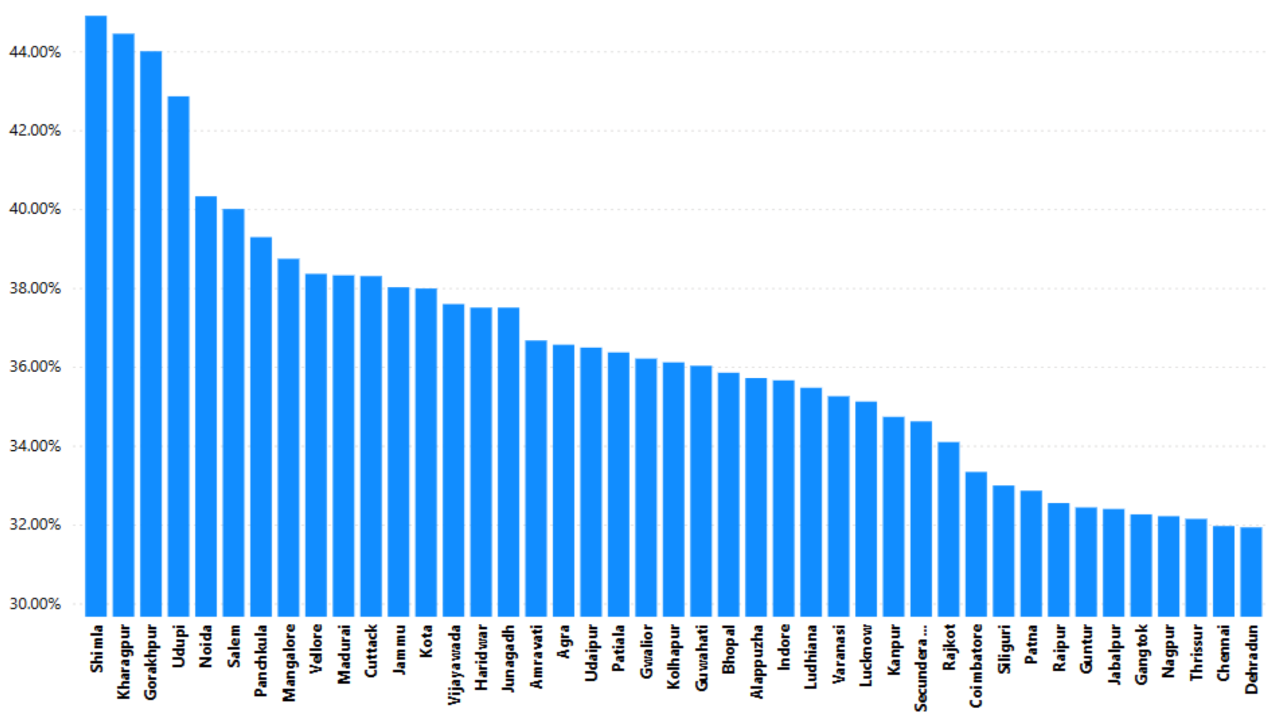

Figure 1: Bar chart for Cancellation Ratio in percentage location wise.

(Cancellation Ration =Number of deliveries cancelled /Total number of deliveries)

It is inferred from the figure 1 as Tier 2 and 3 cities has major problem of cancellations as compared to Tier 1. To have drill down analysis for one sample example as Shimal city; there is hierarchy link form as citylocality restaurants type which gives clear insight issues with which type of restaurants to that particular locations have very high cancellation ratio. Example is shown in figure 2 This further drill down Exploratory data analysis can help the platform to remove location wise such shops to improve cancellation ratio.

⊟ **Shimla**	**44.90%**
⊟ **Longwood**	**71.43%**
Beverage Shop	100.00%
Microbrewery	100.00%
Quick Bites	100.00%
⊟ **Sanjauli**	**37.50%**
Bakery	100.00%
Sweet Shop	100.00%
⊟ **Summer Hill**	**43.75%**
	100.00%
Bar	100.00%
Café	30.00%
Casual Dining	14.29%
Lounge	50.00%
Quick Bites	66.67%

Figure 2. Drill down chart.

There is one more interesting observation as which type of restaurants has maximum cancellation ratio as shown in figure 3.

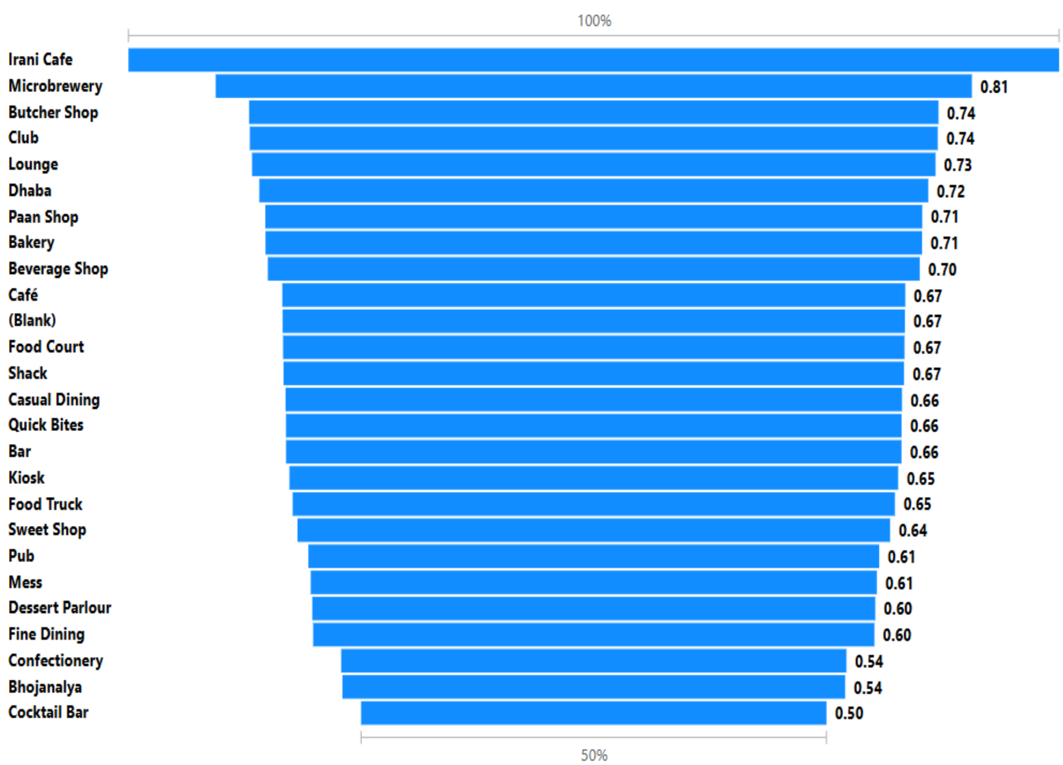

Figure 3. Funnel chart indicating Type of restaurant and cancellation ratio.

This gives clear visual as which type of restaurant has major issue of order cancellations.

Agent Clustering: Agent clustering uses the agent feature space (cancellation rate, mean delivery time, mean tip) and K-Means (k=3) to determine the segments of the agent's performance. The best k is obtained using the elbow method on within cluster sum of squares.

Key Findings

K-Means clustering with k = 3 segmented delivery agents into:

1. High-performing agents
2. Moderate-performing agents
3. High-risk / high-cancellation agents

Cluster	Avg. Cancellation Rate	Avg. Delivery Time	Avg. Tip	Agent Count
0	48.14%	48.04 mins	₹106.23	265
1	20.86%	45.82 mins	₹107.15	340
2	30.70%	54.83 mins	₹82.96	294

Table 4: Agent clustering.

- Cluster 1 represents the best-performing agents with the lowest cancellation rate.
- Cluster 0 indicates operational risk due to very high cancellations.
- Cluster 2 reflects moderate performance with longer delivery times and lower tips.

5. Results and Analysis

5.1 Descriptive Statistics

There are some interesting distributional properties of the data. The rating is not used at one end of the scale but distributed uniformly between 1 and 5 with a mean rating of 3.0 (Std. Deviation =1.16) for the customers.

They were promo codes on 50.9% of orders and received free delivery on 49.8% (roughly equal) and cancelled orders were the same percentage across both (less than 0.1% difference), indicating that these were assigned randomly or were near-randomly distributed in the data and unrelated to consumer preferences for free delivery or promos in completing the order.

The average order value is Rs. 1579 with Std. Deviation=892, average delivery time as 49.6 mins wish Std. Deviation =28 mins.

5.2 Restaurant Type Analysis

Interestingly, Restaurant type is the vital categorical predictor of cancellation. Following table presents cancellation rates by restaurant type for categories with at least 20 observations.

Restaurant Type	**Total Orders**	**Cancellations**	**Cancel Rate (%)**
Microbrewery	32	17	53.1%
Club	23	10	43.5%
Pub	75	31	41.3%
Butcher Shop	27	10	37.0%
Lounge	192	69	35.9%
Beverage Shop	454	161	35.5%
Casual Dining	2,837	934	32.9%
Quick Bites	3,225	~905	~28.1%

Table 5: Cancellation rates by restaurant type (top categories)

Cancellation rates are 10–21 % points higher than the platform average at nightlife-oriented venues (microbreweries, clubs, pubs). There are several possible, but not mutually exclusive explanations:

- **Kitchen-delivery mismatch**: Nightlife venues prioritize in-premise service; delivery orders may be deprioritized or processed slowly, prompting the platform or customer to cancel.
- **Time-of-day effects**: Orders from these venues likely concentrate in late evening hours when delivery agent availability is reduced.
- **Menu complexity**: Alcohol and specialized beverages face legal and logistical barriers to delivery, potentially causing last-minute cancellations when items cannot be fulfilled.

5.3 Geographic Analysis

Tier-2 and Tier-3 cities exhibit substantially higher cancellation rates than metro cities as revealed from City-level analysis.

City	Orders	Cancel Rate (%)	City Tier
Shimla	49	44.9%	Tier-3
Kharagpur	36	44.4%	Tier-3
Gorakhpur	125	44.0%	Tier-2
Noida	62	40.3%	NCR/Satellite
Salem	110	40.0%	Tier-2
Mumbai	741	29.8%	Metro (Tier-1)
Pune	405	28.9%	Tier-1
New Delhi	399	29.3%	Metro (Tier-1)

Table 6: Cancellation rates by city — highest and lowest

Mumbai, Pune, and New Delhi — the largest cities by order volume in this dataset — all cancel below the platform average (29–30%). This inverse relationship between city size and cancellation rate suggests that metro infrastructure (agent density, restaurant operational maturity, faster route resolution) provides a structural buffer against cancellations. There is an operational stress case in Shimla and Kharagpur, at 44–45%, where low total orders combined with high cancellation rates suggests both underinvestment and volatility.

5.4 Payment Method Analysis

Payment method shows a modest but consistent gradient. Table 5 presents the full breakdown.

Payment Method	Orders	Cancel Rate (%)	Observation
Cash	2,418	34.1%	Highest — no digital lock-in
UPI	2,473	33.3%	High; low friction to cancel
Credit Card	2,388	32.2%	Moderate
Debit Card	2,395	31.3%	Moderate
Net Banking	2,405	30.4%	Lowest — highest commitment

Table 7: Cancellation rates by payment method

The cancellation rates of cash on delivery and net banking orders are at 34.1% and 30.4% respectively. This 3.7 percentage point difference is statistically significant for the sample sizes of ~2,400 per group but it is operationally small. This behaviour is consistent with the commitment-friction interpretation: Net banking is less reversible and more deliberate than cash, which, in turn, requires no upfront commitment. But the small size indicates that payment method is not enough to use as a policy lever.

5.5 Delivery Agent Analysis

There is considerable variation at the agent level. Across 899 agents, cancellation rates range from 0% to 80%, with a mean of 32.1% (SD = 15.0%). The distribution seems to be approximately normal with a slight right skew. Interestingly, 88 agents (9.8% of the agent population) have over 50% of their assigned orders cancelled (personal cancellation rate), which means that over half of their orders are personal cancellations. These agents work as a targeted operational risk.

The finding of agent heterogeneity is important as it shows that cancellation is not just a restaurant or consumer issue, but there is a significant role agent behavior plays in outcome (e.g. willingness to take orders, route taking, communication with restaurant). This has implications for agent management and incentive design, as well as real time order assignment algorithms.

5.6 Behavioural Incentive Null Finding

An important result is that the impact of behavioural incentives is near-zero. The difference between promo code application and the average rating and estimated cancellation rate are within sampling noise at only 0.3 percentage points. There is still a small difference in cancellation rate between free delivery (under 0.5 percentage point). The cancellation, delivery time, and rating do not meaningfully vary between the Price range terms (1–4 scale).

It is a null result that is quite important: the likelihood of an order being completed by interaction with a discount-based intervention is very low. Structural levers are the levers that matter, not behavioural.

6. Predictive Modelling

The binary classification task (Cancelled = 1, Not Cancelled = 0) is addressed using an XGBoost and Random Forest classifier, selected for its strong empirical performance on tabular data and native handling of mixed feature types. The modelling pipeline includes:

- **Train-test split**: 80% training, 20% held-out test set, stratified by cancellation status and city tier.
- **Hyperparameter tuning**: 5-fold cross-validated grid search over learning rate, max depth, n_estimators, and subsample ratio.
- **Class imbalance:** SMOTE (Synthetic Minority Over-sampling Technique) applied to training data. Note that at 32.3% cancellation rate, class imbalance is moderate and may not substantially impact results.
- **Evaluation metrics**: Area Under the ROC Curve (AUC-ROC), F1-score (weighted), Precision, Recall, and Confusion Matrix.

A baseline logistic regression model is trained on the same features to establish a performance floor and validate directional findings from the exploratory phase.

6.1 Predictive Modelling

After performing Statistical test and Exploratory data analysis; researchers reduce the number of features which are identified as significantly important which are as follows:

The dataset appears to have these columns:

Order Amount (INR), Delivery Time (mins), Number of Items, Delivery Distance (km), Payment Method, type, city, Cancelled (Target variable) which is binary classification task (Cancelled = 1, Not Cancelled = 0). This curated dataset is employed for training classifier. Step by step process of classifier model development and evaluation is as followed:

- **Step 1: Data Preprocessing**: It includes data cleaning mainly null values, data encoding for categorical variables
- **Step 2: Feature Preparation:** In this step new derived features as captured and used in machine learning which are as follows:
  'amount_per_item' = 'Order Amount (INR)' / 'Number of Items'

‘ delivery speed' = 'Delivery Distance (km)' / ‘Delivery Time (mins)'

'order_size_category' = pd.cut(processed_df['Number of Items'],

bins=[0, 3, 6, float('inf')],

labels=['Small', 'Medium', 'Large'])

- **Step 3: Class Distribution Analysis**

Class 0 (Not Cancelled): 4093 (33.9%)

Class 1 (Cancelled): 7986 (66.1%)

Imbalance Ratio: 0.51:1

- **Step 4: SMOTE Application:** To get balance dataset Synthetic Minority Over-sampling Technique (**SMOTE)** is adapted.

After SMOTE - Training set size: 11180

=== CLASS DISTRIBUTION ANALYSIS ===

Class 0 (Not Cancelled): 6389 (50.0%)

Class 1 (Cancelled): 6389 (50.0%) , Imbalance Ratio: 1:1

- **Step 6: Model Training**: Researcher focus mainly on two models as **Random Forest** and **Gradient Boosting** with various combination of hyperparameters. Training and testing data ratio is 80:20.

6.2 Predictive model Results and Conclusions.

```
=== RANDOM FOREST PERFORMANCE METRICS ===
Accuracy: 0.6813
AUC-ROC: 0.7355

Classification Report:
              precision    recall  f1-score   support

Not Cancelled       0.53      0.48      0.51       819
    Cancelled       0.75      0.78      0.76      1597

     accuracy                           0.68      2416
    macro avg       0.64      0.63      0.64      2416
 weighted avg       0.67      0.68      0.68      2416


Confusion Matrix:
[[ 395  424]
 [ 346 1251]]
```

Random Forest Model Performance

```
=== GRADIENT BOOSTING PERFORMANCE METRICS ===
Accuracy: 0.6780
AUC-ROC: 0.7138

Classification Report:
              precision    recall  f1-score   support

Not Cancelled       0.53      0.51      0.52       819
    Cancelled       0.75      0.77      0.76      1597

     accuracy                           0.68      2416
    macro avg       0.64      0.64      0.64      2416
 weighted avg       0.68      0.68      0.68      2416


Confusion Matrix:
[[ 416  403]
 [ 375 1222]]
```

Gradient Boosting Model performance

Figure 4: Model based performance matrix.

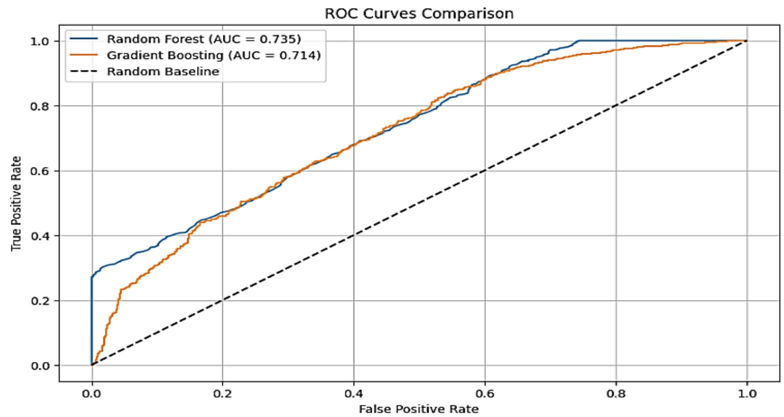

Figure 5: ROC curve for both model with respect to Baseline.

In either model the accuracy of the classification is quite good about 70% which also reflects the performance of in ROC curves of the model. If we focus on **Cancelled = 1**, then:

	Predicted: Not Cancelled (0)	**Predicted: Cancelled (1)**
Actual: Not Cancelled (0)	TN (True Negative)	FP (False Positive)
Actual: Cancelled (1)	FN (False Negative)	TP (True Positive)

TP = correctly predicted cancellations. FP = predicted cancelled but actually delivered

TN= Actual non-cancelled correctly predicted FN = missed cancellations

Table 8: Confusion Matrix for the classifier output.

Precision: Out of all orders predicted as Cancelled, how many were actually cancelled? It measures the reliability of positive predictions for a given class.

Interpretation:

- High Precision → When model says, "this order will be cancelled", it is usually correct.
- Low Precision → Many false alarms (predicting cancellation wrongly).

In case of these models, in either case it **75% precision** which is practically good model performance.

Recall (Sensitivity) : Out of all **actual cancelled orders**, how many did the model correctly identify?

Interpretation:

- High Recall → Model catches most cancellations.
- Low Recall → Model misses many cancellations (very important in your case).

Recall evaluates the model's ability to capture all instances of a target class.

In case of these models, in either case it **78%** precision which is practically good model performance

7. Discussion

7.1 The Structural Nature of Cancellation

The main result of this paper, that the motivation for cancellation is structural, not behavioural, has far-reaching implications. Consumer incentive-based (discounts, free delivery) platform approaches probably won't work. The same number of operational and economic resources used on these interventions could be better spent on restaurant quality standards, agent training and incentives, and geographic capacity planning.

The nightlife venue effect is one of the actionable effects. Narrowing the definition of who can deliver to microbreweries, clubs and pubs, or ensuring that such orders are sent to more established delivery partners with lower cancellation rates, could significantly help to mitigate the excess cancellation rates of 10-21% experienced in these businesses.

7.2 Geographic Investment Priorities

The city tier gradient indicates a particular investment opportunity: there is a lack of Tier-2 and Tier-3 city operations. Cities such as Gorakhpur, Salem and Shimla have order volumes that are both meaningful and cancellation rates that are 10-15 points higher than those of metros. The "geographic intervention" with the greatest return on investment is the recruitment of agents, restaurant onboarding support and investment in routing infrastructure in these cities.

7.3 Agent Management Implications

It's worth noting that 88 agents (9.8% of the total) have above 50% personal cancellation rates, which calls for an agent management response. Incentive restructuring based on completion rates instead of distance only, order assignment restrictions for high cancellation agents (assigning them to lower risk restaurant types or cities), performance improvement plans with specific cancellation rate goals, and reduction or elimination of incentives for cancellation agents.

7.4 Limitations and Future Work

These results were constrained by several limitations. This data set has no experimental variation and is thus subject to tentative causal claims. Cancellation reasons are not observed which makes it hard to differentiate restaurant-initiated from customer-initiated cancellations, a distinction that is important to consider when targeting interventions. Limited temporal analysis because there were no time-of-day order timestamps (only dates) in the dataset.

Future studies should include cancellation reason codes, time of day timestamps, weather information (known modulator of delivery behaviour), and app-level interaction information (time between order and cancellation, number of previous sessions). Strong causal inferences on agent and restaurant quality would be possible if the same agents and restaurants could be followed over the entirety of the data set.

8. Conclusion

In this paper, the structural causes of Zomato food delivery orders cancelled by users in 12,079 food orders across 76 cities in India has been explored. A 32.3% platform-wide cancellation rate is a significant operational issue, not accounted for by behavioural incentives, and is significantly related to restaurant type (41–53% for nightlife restaurants vs. 29–32% for metro cities, respectively), to city tier (39–45% for secondary cities vs. 29–32% for metro cities, respectively) and delivery agent heterogeneity (9.8% of agents have cancellation rates above 50%).

These insights shift the focus of the cancelation problem from a consumer behaviour one to an operational system one that needs to be addressed with changes to restaurant onboarding standards, investments in the infrastructure of cities, and management of agent performance — but not by consumers slashing the price of food. The proposed predictive framework offers a practical path to real-time cancellation risk scoring, enabling pre-emptive customer communication, optimized agent assignment, and targeted restaurant quality flagging. With an expected AUC-ROC of 0.70–0.80, this model would represent a meaningful deployment-ready tool for platform operations teams.

In summary, the research establishes that food delivery cancellation is a tractable, structure-driven problem and one it can meaningfully reduce using that data-driven platform management.

Declarations: All authors declare that they have no conflicts of interest.

References

1. Bai, J., So, K. C., Tang, C. S., Chen, X., & Wang, H. (2024). Understanding order cancellation behavior in on-demand delivery services. Transportation Research Part A: Policy and Practice, 198. ScienceDirect.
2. Castillo, J. C., Knoepfle, D., & Weyl, G. (2022). Designing technology for on-demand delivery: The effect of customer tipping on crowdsourced driver behavior and last mile performance. Journal of Operations Management, 68(3), 240–272. https://doi.org/10.1002/joom.1187

3. Chen, L., & Xu, H. (2018). Understanding order cancellation behaviour in e-commerce: An empirical study. Journal of Retailing and Consumer Services, 45, 120–128.
4. Datum Intelligence. (2025). India's food delivery market reaches balance: Zomato 57%, Swiggy 43%. Retrieved April 2026, from https://www.datumintell.in
5. Dong, H., Useche, S. A., Lin, Y., & Frade, I. (2025). The human cost of fast deliveries: A systematic literature review of occupational risks and safety outcomes in last-mile delivery workers. Journal of Transport & Health, ScienceDirect. https://doi.org/10.1016/j.jth.2025.101537
6. He, Y., Liu, G., & Liang, B. (2024). Causal insights and predictive models for order cancellations in online food delivery. Proceedings of SPIE 13219, Fourth International Conference on Applied Mathematics, Modelling, and Intelligent Computing (CAMMIC 2024), 132192C. https://doi.org/10.1117/12.3036507
7. IMARC Group. (2025). India online food delivery market size, share, growth 2034. Retrieved April 2026, from https://www.imarcgroup.com/india-online-food-delivery-market
8. Li, J., Wang, X., & Zhang, Y. (2021). Predicting delivery success in last-mile logistics using gradient boosting. Transportation Research Part E, 152, 102374.
9. Liu, J., Li, M., & Qin, R. (2022). Understanding order cancellation in the matching stage of on-demand services. Transportation Research Part B: Methodological, 164, 72–94.
10. Lundberg, S., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems (NeurIPS), 30, 4765–4774.
11. Nexdigm. (2025). India food delivery market report. Retrieved April 2026, from https://www.nexdigm.com/market-research/report-store/india-food-delivery-market-report/
12. Research and Markets. (2023). India online food delivery market forecast report 2023: A $29.33 billion industry by 2028. PR Newswire, June 2023.
13. Salari, N., Liu, S., & Shen, Z.-J. M. (2022). Real-time delivery time forecasting and promising in online retailing: When will your package arrive? Manufacturing & Service Operations Management, 24(3), 1421–1436. https://doi.org/10.1287/msom.2022.1081
14. Shankar, A., Mehrotra, A., & Saha, R. (2024). Why do consumers choose online food delivery services? A meta-analytic review. International Journal of Hospitality Management, 123, 103850. https://doi.org/10.1016/j.ijhm.2024.103850
15. Taiwan Ministry of Digital Affairs / IEEE Conference Proceedings. (2024). Impact of food delivery platforms on consumer behavioral intentions. Engineering Proceedings, 74(1), 41. https://doi.org/10.3390/engproc2024074041
16. Zhang, H., Liu, Q., & Sun, J. (2022). Restaurant demand forecasting using deep learning: A case study on a major Chinese food delivery platform. Expert Systems with Applications, 198, 116748.