



Data-Driven Strategy for Merchant Incentive Optimization in Digital Payment Ecosystems

Hira Ajmal

Lahore School of Economics, Lahore, Pakistan

Article
Open Access
Published

ABSTRACT

In the rapidly evolving digital payment ecosystem, optimizing merchant incentives has become a strategic necessity for increasing transaction volume, improving platform loyalty, and fostering sustainable customer engagement. Traditional incentive schemes often rely on heuristic or fixed-rate models that fail to account for the dynamic and heterogeneous nature of merchant behavior. To address these limitations, this study introduces a novel data-driven framework that employs graph-based representation learning to model and analyze complex transactional interdependencies among merchants and consumers. By capturing structural and behavioral similarities through graph neural networks (GNNs), the proposed approach enables precise prediction of each merchant's sensitivity to diverse incentive strategies. The model integrates transaction frequency, geographical proximity, and customer overlap to generate interpretable embeddings that guide optimized budget distribution under financial constraints. Through an integrated optimization layer, incentives are allocated based on predicted responsiveness, ensuring that marketing expenditures are directed toward high-impact segments. Real-world experiments conducted on large-scale digital payment datasets validate the framework's effectiveness, demonstrating substantial improvements in merchant participation rates, transaction growth, and cost efficiency compared to baseline regression and deep learning models. The findings highlight the potential of combining machine learning, network science, and marketing analytics to design more adaptive and data-intelligent promotional systems, thereby paving the way for scalable and targeted incentive management in next-generation digital financial platforms.

Keywords

- Graph Neural Networks
- Dynamic Pricing
- Merchant Incentives
- Digital Payments
- Representation Learning

Received: May 02, 2025; Revised: October 28, 2025; Accepted: October 28, 2025; Published: November 01, 2025

© 2025 REPA. All rights reserved.

1. Introduction

The proliferation of digital transaction platforms, including Alipay and Apple Pay, has reshaped financial interactions in contemporary society [1]. A key goal for mobile payment providers is to expand their user base, fostering a shift from conventional payment methods to digital alternatives for both businesses and consumers.

To drive adoption, service operators deploy various promotional tactics, primarily through merchant-based incentives. These incentives, encompassing commission structures, discount vouchers, and supplementary business services, are tailored to encourage vendors to process payments via designated digital platforms [2]. For example, Alipay assigns a distinct incentive QR code to each vendor, enabling them to extend rewards to customers. These rewards can be availed solely through Alipay transactions, and in exchange, merchants accumulate commissions based on the volume of successfully redeemed incentives [3].

This framework, in which benefits are initially distributed to customers and subsequently reclaimed through transaction completions, establishes a model where higher payment activity directly correlates with increased merchant revenue [4].

Historically, commission structures for merchants have relied on predefined allocation strategies. However, merchant engagement with incentives varies considerably. Certain vendors demonstrate a strong dependency

between commission rates and transaction frequencies, whereas others exhibit minimal responsiveness to fluctuating incentive levels. Accurately identifying merchants with pronounced sensitivity to promotional offers is essential for optimizing marketing expenditure within predefined financial limits [5].

Modeling the dependency between promotional incentives and transaction frequency introduces several complexities. The prolonged impact of incentives on merchant engagement must be accounted for, as vendors may only perceive substantial benefits after prolonged participation [2]. Additionally, financial constraints restrict the feasibility of conducting exhaustive empirical evaluations to comprehensively define the incentive-response function for every merchant [4].

Instead, an experimentally constrained methodology is frequently utilized, where predetermined incentive values, sampled from a fixed distribution (e.g., 1, 2, 5, 10, 20 units of currency), are randomly assigned, followed by an observation of transactional responses [3].

Two principal obstacles emerge in this setting. Firstly, an effective representation learning approach is necessary to cluster merchants based on their responsiveness to incentives, allowing for robust statistical inference of incentive-response relationships within these groups [6, 7]. Secondly, the incentive transaction dependency for individual vendors can exhibit intricate patterns,



necessitating the incorporation of informed prior knowledge to mitigate estimation variability [8].

This research introduces a graph-based learning framework to enhance the optimization of merchant rewards in digital payment marketing. Vendor behavior is analyzed through customer purchasing trends, extracting meaningful insights from geographical transaction distributions. The evaluation of incentive-response relationships, derived from real transactional datasets, reveals that transaction frequencies demonstrate a linear and consistent increase with rising incentive levels. Experimental validation substantiates the efficacy of this methodology, showcasing enhanced marketing performance. Finally, the problem of distributing incentives under financial limitations is formulated as an optimization problem [5], leveraging a linear programming approach [3], with empirical testing conducted through real-world deployment on the Alipay platform.

2. Background

This section presents a concise review of relevant literature, emphasizing two core areas: dynamic pricing methodologies and graph-based feature learning. These principles serve as the basis for the proposed strategy.

2.1. Dynamic pricing strategies

Dynamic pricing utilizes analytical techniques to tackle two fundamental issues: (1) forecasting consumer reactions to diverse pricing schemes [2,3]; and (2) formulating optimal price points to fulfill business objectives [4,5]. These techniques allow businesses to evaluate promotional effectiveness and identify strategic price settings to enhance revenue within specified timeframes.

The framework of dynamic pricing is generally composed of two key phases. Initially, predictive models leveraging machine learning algorithms are applied to estimate product demand [8], infer purchasing tendencies from incomplete transaction datasets, or analyze historical sales data to anticipate future market behavior. Subsequently, these estimations are integrated into optimization models designed to establish pricing mechanisms that align with corporate objectives [5,9].

In the domain of mobile financial transactions, service providers adopt an analogous approach, striving to expand user engagement by incentivizing specific payment platforms [2]. Conventional studies predominantly rely on linear regression models to capture sales trends [4]. However, evaluating merchant receptivity to varying incentive structures introduces a distinctive challenge [5].

Accurately determining incentive responsiveness across a vast network of merchants, particularly under constraints of limited data availability, remains an under-explored subject. This research proposes a novel mechanism for quantifying merchant-specific incentive adaptability, extending beyond traditional dynamic pricing methodologies.

2.2. Graph-based representation learning

Graph-based representation learning has gained substantial attention, particularly with the advent of graph neural networks (GNNs) [6,7,10], which encapsulate subgraph structures into latent embeddings. Consider an undirected graph $G = (V, E)$ where V represents a set of N nodes, and E consists of $|E|$ edges (i, j) . The graph structure is defined by a sparse adjacency matrix $A \in \mathbb{R}^{N \times N}$, a node attribute matrix $X \in \mathbb{R}^{N \times P}$, and a graph Laplacian operator formulated as $L = I - D^{-1/2}AD^{-1/2}$. These techniques utilize local neighborhood aggregation to construct informative node embeddings.

An influential framework, GraphSAGE, introduced a mechanism for aggregating neighboring node information via different pooling operations [6]. The general update rule can be described as:

$$H^{(t+1)} = \sigma(\text{CONCAT}(\psi(A, H^{(t)}), H^{(t)})W^{(t)}) \quad (1)$$

where $H^{(t)} \in \mathbb{R}^{N \times K}$ denotes the hidden layer representation at depth t , initially set as $H^{(0)} = X$. The transformation matrix $W^{(t)}$ is trained for each layer, σ represents a non-linear activation function, and $\psi(\cdot)$ is a neighborhood aggregation function, such as mean or max pooling. By stacking multiple layers, the model assimilates information from broader graph neighborhoods, enhancing node embeddings.

To improve adaptability, attention-based strategies have been introduced, enabling dynamic neighborhood weighting based on relevance. Additionally, path-based filtering methods have been explored to refine aggregation regions, allowing models to learn more discriminative node representations. These methodologies have demonstrated superior performance across various benchmark datasets, making them instrumental in diverse applications.

Within the domain of mobile transaction networks, GNNs are employed to encode merchant characteristics based on consumer purchasing patterns [11, 12]. This technique helps uncover variations in customer engagement, yielding structured embeddings that effectively describe merchant behaviors. Merchants with a wider customer reach generally exhibit stronger transaction activity, while those catering to niche demographics often display lower responsiveness to promotional incentives. Furthermore, spatial proximity frequently results in customer overlap, leading to comparable merchant incentive adoption patterns.

By structuring transactional relationships through graph-based modeling, an optimized strategy for incentive allocation can be systematically developed, as explored in subsequent sections.

2.3. Online experimentation and sensitivity analysis

To evaluate the influence of financial incentives on merchant engagement, a controlled online study was carried out, wherein varying incentive magnitudes were

randomly allocated [5]. This methodology facilitated the derivation of the objective-incentive function, capturing how business metrics evolve in response to incentive modifications. Data collection was performed across several days, incorporating a diverse range of merchants transacting on the platform [3].

Given that each merchant's reaction to incentives is observable under only one experimental condition at any given moment, extrapolating these reactions to merchants with similar characteristics becomes essential [9]. It is postulated that merchants exhibiting comparable transactional patterns display analogous responsiveness to financial stimuli. Regional trends in sensitivity are depicted in

Figure 1, highlighting that neighboring regions often exhibit correlated behaviors. This observation underscores the utility of graph-based learning models in capturing such structural relationships.




Figure 1. Regional variations in incentive responsiveness. The x -axis and y -axis correspond to longitude and latitude, respectively. Higher values indicate increased sensitivity to financial incentives.

2.4. Graph neural network-based framework

Graph neural networks (GNNs) are employed to model interdependencies among merchants by constructing a transaction-based network [6,7]. This network is defined as $G = (V, E)$, where V comprises merchants and customers, while E denotes financial transactions. An adjacency matrix $A \in \{0,1\}^{N \times N}$ is formulated, where $A_{i,j} = 1$ signifies an interaction between nodes i and j [13].

Merchant-specific features $X \in \mathbb{R}^{N \times P}$ encode business attributes, while transaction-related attributes $Z \in \mathbb{R}^{|E| \times D}$ encapsulate details of financial exchanges [9]. The GNN propagates information iteratively through the following update mechanism:

$$h_i^{(t+1)} = \sigma \left(W^{(t)T} \psi \left(\{h_j^{(t)} \mid j \in N(i) \cup i\}; \theta \right) \right) \quad (2)$$

where $h_i^{(t)}$ represents the feature embedding at layer t , σ is a non-linear activation function, and $\psi(\cdot)$ is an aggregation operator parameterized by θ . Initial feature embeddings are initialized as follows:

$$h_i^{(0)} = W_x^T X_i + \sum_{j \in N(i)} W_e^T Z_{(i,j)} \quad (3)$$

2.5. Monotonic mapping and optimization strategy

The learned merchant embeddings from the GNN are mapped to an objective-incentive function through a constrained transformation [10]:

$$f(i, c) = c \cdot \text{SOFTPLUS}(W_g^T h_i^{(T)}) + \text{RELU}(W_p^T h_i^{(T)}) \quad (4)$$

where W_g and W_p project the embeddings into parameters defining the gradient and intercept of the mapping function. This ensures a monotonic relationship between incentives and business performance, where the gradient g_i characterizes responsiveness [14].

To determine optimal parameters $\beta = \{W, W_x, W_e, \theta, W_g, W_p\}$, the following loss function is minimized:

$$\min_{\beta} \sum_{i,c} \mathcal{L}(f(i, c; \beta), y_{i,c}) \quad (5)$$

where $\mathcal{L}(\cdot)$ denotes the mean absolute error between predicted and actual business performance outcomes $y_{i,c}$. Model training is executed using the ADAM optimization algorithm with mini-batch updates in TensorFlow [15,16].




Figure 2. Graph neural network architecture integrating monotonic mapping for incentive impact estimation.

2.6. dataset overview

The characteristics of the datasets employed in the experimentation are outlined in Table 1 [2,4].

Table 1: Overview of datasets used for experimentation.

Dataset	V	E	Node Attributes	Edge Attributes	Labeled Samples
D1	90.08×10^6	161.7×10^6	4998	86	2.18×10^6
D2	97.18×10^6	172.1×10^6	4998	86	2.31×10^6

2.7. Optimized incentive allocation via linear programming

The allocation of incentives is structured as a linear programming task, utilizing the estimated function $g(i, c)$ [3]. The objective is to allocate the most beneficial intervention to each vendor while adhering to a budgetary limitation [4].

$$\max \sum_{i,c} g(i, c) b^*(i, c) \quad (6)$$

subject to:

$$\sum_{i,c} b^*(i, c) \cdot c \leq B \quad (7)$$

where the best intervention m satisfies:

$$m = \arg \max_c g(i, c) \cdot c \cdot \mu \quad (8)$$

with μ representing the optimal dual variable. The effectiveness of this methodology is validated within Alipay's digital framework, with findings detailed in the next segment.

3. Experimental evaluation

This segment provides a comprehensive analysis of the system's efficiency. First, offline performance is scrutinized using datasets derived from digital trials. Subsequently, online evaluation is conducted by implementing the proposed framework within Alipay and comparing it against a multi-layer perceptron model that adopts the optimal strategy elaborated in Section 3.3. Due to confidentiality constraints, certain numerical specifics that might be deemed sensitive are not revealed.

3.1. Experimental configuration

This section elaborates on the dataset employed for training the model and the corresponding experimental settings [2].

Data Acquisition: The datasets originate from two distinct digital trials conducted over a span of 15 days each. Data gathering involved the arbitrary selection of 0.2% of vendors, who were then categorized into different test groups, each subjected to a predefined intervention. In total, 13 experimental groups were designed, each corresponding to a distinct treatment to maintain an impartial allocation of merchants based on their responsiveness to promotional strategies [15].

The datasets encompass over two million tagged vendors, with recorded business performance indicators, such as the volume of transactions finalized within the subsequent three days and the number of active days with completed sales in the same duration. The model is structured on a transactional network, where links symbolize interactions between labeled merchants and their two-hop vicinity. This incorporates direct exchanges between merchants and customers, as well as secondary interactions involving clients and additional vendors. The resultant transactional network comprises over 90 million entities, encompassing merchants and consumers, with

hundreds of millions of transactional connections. A summary of the datasets utilized in offline trials is depicted in Table 2. Data from two separate timeframes were examined to validate the consistency of the observations.

Table 2: Overview of experimental datasets.

Dataset	Number of Vendors	Number of Transactions
1	X	Y
2	X	Y

Benchmarking Approaches: To evaluate the efficacy of the proposed method, comparisons were made against conventional regression techniques. Linear models, neural networks, and ensemble-based models are widely utilized for regression tasks. However, given the sparse nature of the generated attributes, ensemble methods were omitted as they are less effective in handling high-dimensional sparse data. Instead, the performance was assessed using linear regression (LR) and a deep learning model (DNN) featuring a multi-layer perceptron (MLP) structure with a constrained monotonic transformation in the final layer, as discussed in Section 3.2.2. The suggested framework is denoted as the GE model throughout the experimental assessments.

The architecture of the DNN model consists of two layers, each employing an embedding dimension of 256. Similarly, the graph-based model maintains a two-layer depth, allowing each labeled merchant to integrate information from direct consumers (one-hop connections) and merchants sharing common customers (two-hop connections). The embedding size remains fixed at 256. Key hyperparameters, including the learning rate and penalty terms, were fine-tuned using a grid search approach.

A stratified random sampling technique was employed, with 80% of the merchant data allocated for model training and the remaining 20% utilized for validation. The models were evaluated based on standard regression performance indicators, with the results displayed in Table 3.

Table 3: Comparison of Mean Absolute Error (MAE) between regression models.

Dataset	Model	MAE
1	LR	0.1432
	DNN	0.1404
	GE	0.1357
2	LR	0.1441
	DNN	0.1409
	GE	0.1361

4. Predictive modeling evaluation

Empirical validation was performed on real-world transaction data obtained from controlled online trials. To measure predictive accuracy, Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) were employed as standard performance metrics. While these indicators do not directly quantify the model's effectiveness in

identifying individual merchant responsiveness, they provide insights into its overall predictive reliability across diverse treatment strategies.

Given table presents the evaluation findings. The deep learning model (DNN) demonstrates superior MAE performance relative to linear regression but exhibits a slightly elevated RMSE, implying higher variance. The incorporation of a constrained monotonic mapping layer within the DNN model contributes to mitigating overfitting while improving generalizability. However, its capacity to capture subtle variations across merchants remains constrained. The graph-based embedding (GE) model outperforms both benchmarks in terms of MAE and RMSE, reinforcing its strength in capturing transaction patterns and modeling incentive responses.

5. Analysis of sensitivity to incentives

Further investigation was conducted on the estimated gradient, to understand how merchants react to varying incentive levels [5]. This gradient acts as a proxy for merchant sensitivity to marketing strategies, delineating the response-incentive curve. Ensuring precise estimation of these gradients remains a crucial objective.

Unlike direct numerical validation, the accuracy of gradient estimations was examined through comparative assessment [4]. A well-trained model should effectively differentiate between merchants with distinct sensitivities, wherein those categorized as highly responsive exhibit greater improvements in commercial activity under enhanced incentive treatments.

Merchant transactional performance under differing incentive allocations is denoted as y_h (higher incentive) and y_l (lower incentive), respectively. The improvement in performance, termed as uplift gain, is computed as $u = y_h - y_l$. Merchants classified as more responsive to incentives should display higher uplift values relative to those categorized as less sensitive [6].

To validate this approach, the test dataset was utilized to infer gradients per merchant, followed by sorting in descending order. This enabled the segmentation of merchants into two groups: highly sensitive and less responsive categories, designated as Ω^+ and Ω^- , respectively. The model's effectiveness in optimizing incentives was assessed based on its ability to accurately differentiate between these groups.

For the highly responsive category (Ω^+), the uplift gain was determined as $u^+ = y_h^+ - y_l^+$, whereas for the less sensitive category (Ω^-), the corresponding uplift was denoted as u^- . An ideal model would exhibit a significant difference between u^+ and u^- . As depicted in Figure 3, the GE model yields a markedly greater uplift difference compared to the DNN model, underscoring its superiority in capturing variations in merchant responsiveness [9].

Since linear regression does not generate personalized gradients, its results were omitted from this analysis. For both DNN and GE models, merchants in the test set were ranked based on inferred gradients and divided into five equally sized cohorts. The GE model consistently surpasses the DNN model for the most responsive merchants [17]. However, for the least sensitive cohort, the GE model exhibits a flatter response-incentive curve, suggesting a refined understanding of merchants with minimal reaction to incentives.

6. Digital platform deployment results

The developed model was implemented within a mobile transaction ecosystem using an industry-standard A/B testing methodology. Initially, the test group encompassed 1% of merchants, with data being recorded for five consecutive days. The exposure rate was then progressively increased to 2.5%, 5%, and 15%, maintaining the same observation period for each stage. Ultimately, from January 10th to January 14th, A/B testing scaled up to 30% of traffic, serving as the conclusive evaluation phase before broader implementation.

During the experiment, millions of users engaged with the platform under the 30% traffic condition. The effectiveness of the marketing approach was assessed based on expenditure optimization and two key performance indicators: Metric 1, which quantifies the average payment frequency per merchant, and Metric 2, measuring the average number of days with at least one transaction per merchant.

Table 4 showcases the percentage-based comparative improvements of the graph-embedded (GE) model over the deep neural network (DNN) model at the 30% traffic threshold. A confidence interval of 95% is provided, demonstrating a 2.71% decrease in promotional expenditures while yielding statistically significant enhancements in the business performance indicators, with p -values confirming high confidence in the observed effects.

Table 4: Performance improvement (%) of GE model over DNN model at 30% traffic exposure.

Model	Cost Reduction (%)	Metric 1 (%)	Metric 2 (%)
Baseline	-	-	-
GE Model	-2.71% [-3.06%, -2.36%]	[+0.28%, 0.52%]	[+0.29%, 0.54%]

Figure 3 illustrates the observed trajectory of relative performance gains across the two business objectives, following the adoption of the GE model's optimization strategy. The rising trend suggests a steady improvement in overall transactional outcomes, potentially attributed to merchants adapting their behavior in response to incentive-driven marketing initiatives.




Figure 3. Relative performance enhancement on business objectives (30% traffic exposure).

7. Conclusion

This research introduces a refined approach to tackling marketing optimization challenges within large-scale mobile financial ecosystems. To the best of available knowledge, this represents the pioneering application of graph-based neural learning within a real-world commercial marketing infrastructure. The framework prioritizes the identification of merchants exhibiting heightened receptiveness to promotional interventions by leveraging transaction network structures through advanced representation learning methodologies.

To counteract fluctuations induced by limited exposure groups, a structured linear mapping function has been incorporated, ensuring robustness in feature extraction. Additionally, uplift gains have been established as an innovative metric for model assessment, facilitating optimized decision-making and enhanced resource allocation by minimizing inefficiencies associated with traditional marketing strategies.

Comprehensive evaluation was carried out using an A/B testing framework spanning multiple weeks. The empirical findings consistently validate the superiority of the proposed methodology in boosting marketing efficiency and amplifying commercial outcomes within mobile payment ecosystems.

References

- [1] Clow KE (2004) *Integrated Advertising, Promotion, and Marketing Communications* (Pearson Education India)
- [2] Ferreira KJ, Lee BHA, Simchi-Levi D (2015) "Analytics for an online retailer: demand forecasting and price optimization" *Manufacturing & Service Operations Management* (18(1), pp. 69–88)
- [3] Gallego G, Wang R (2014) "Multiproduct price optimization and competition under the nested logit model with product-differentiated price sensitivities" *Operations Research* (62(2), pp. 450–461)
- [4] Schlosser R, Boissier M (2018) "Dynamic pricing under competition on online marketplaces: a data-driven approach" *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 705–714)
- [5] Ito S, Fujimaki R (2017) "Optimization beyond prediction: prescriptive price optimization" *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1833–1841)
- [6] Hamilton WL, Ying R, Leskovec J (2017) "Inductive representation learning on large graphs" *arXiv preprint arXiv:1706.02216*
- [7] Hamilton WL, Ying R, Leskovec J (2017) "Representation learning on graphs: methods and applications" *arXiv preprint arXiv:1709.05584*
- [8] Chen T, Guestrin C (2016) "XGBoost: a scalable tree boosting system" *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794)
- [9] Schmidhuber J (2015) "Deep learning in neural networks: an overview" *Neural Networks* (61, pp. 85–117)
- [10] Tibshirani RJ et al. (2014) "Adaptive piecewise polynomial estimation via trend filtering" *Annals of Statistics* (42(1), pp. 285–323)
- [11] Chung FRK (1997) *Spectral Graph Theory* (American Mathematical Society)
- [12] Veličković P et al. (2017) "Graph attention networks" *arXiv preprint arXiv:1710.10903*
- [13] Liu Z et al. (2018) "GeniePath: graph neural networks with adaptive receptive paths" *arXiv preprint arXiv:1802.00910*
- [14] Wu Z et al. (2019) "A comprehensive survey on graph neural networks" *arXiv preprint arXiv:1901.00596*
- [15] Kingma D, Ba J (2014) "Adam: a method for stochastic optimization" *arXiv preprint arXiv:1412.6980*
- [16] Abadi M et al. (2016) "TensorFlow: a system for large-scale machine learning" *arXiv preprint arXiv:1605.08695*
- [17] Hofmann T, Schölkopf B, Smola AJ (2008) "Kernel methods in machine learning" *Annals of Statistics* (36, pp. 1171–1220)