IAES International Journal of Artificial Intelligence (IJ-AI)
Vol. 14, No. 6, December 2025, pp. 4848~4864
ISSN: 2252-8938, DOI: 10.11591/ijai.v14.i6.pp4848-4864 ❐ 4848

A merchant analytics framework for revenue forecasting and financial stress detection using transaction data

Yara Harb, Wissam Baaklini, Nadine Abbas, Seifedine Kadry
Department of Computer Science and Mathematics, Lebanese American University, Beirut, Lebanon

Article Info

Article history:

Received Jun 4, 2025
Revised Oct 30, 2025
Accepted Nov 8, 2025

Keywords:

Financial stress detection
Merchant analytics
Prophet model
Time series forecasting
Transaction data

ABSTRACT

By processing payments and providing specialized financial services, acquiring banks are essential for merchants' operations. To forecast 30-day revenue trajectories, identify seasonal demand patterns, and identify early indicators of financial stress, this paper presents a scalable merchant analytics framework that benefits from transactional data. The framework captures multi-level seasonalities using Prophet time series model, allowing dynamic product offerings like revenue-based loans. Proactive risk management is supported offerings like revenue-based loans. Proactive risk management is supported. by a new stress-flagging mechanism that identifies merchants at risk based on deviations in revenue trends. The framework achieved a median 30-day mean absolute percentage error (MAPE) of 56.51% after the validation on a datasetwith 130,350 transactions from 460 merchants in a volatile economic environment. The model demonstrated significant practical utility in identifying financial distress and segmenting merchant behavior, despite its moderate predictive precision, which is common challenge in high-variance merchant datasets. Model outputs are converted into decision-support visualizations along with an interactive dashboard.

Corresponding Author:

Nadine Abbas
Department of Computer Science and Mathematics, Lebanese American University
Beirut, Lebanon
Email: nadine.abbas@lau.edu.lb

1. INTRODUCTION

Acquiring banks offer value-added services such as credit and loans, also, they work on facilitating merchant transactions in the dynamic financial services industry. Offering specialized support, expecting demand spikes, and reducing risks all depend on an understanding of merchant performance as determined by revenue forecasting and financial health evaluation. Since small and medium-sized businesses, which include many retailers, account for around 50% of the world's gross domestic product (GDP), their financial stability is essential to economic expansion [1]. Advanced analytics solutions are necessary because existing approaches frequently miss the complex seasonal patterns and unexpected distress signals present in transaction data.

This paper proposes a merchant analytics framework for revenue prediction, seasonal trend detection, and financial stress identification. The framework forecasts 30-day revenue by using the Prophet time series forecasting model, which has proven strong performance effectively in financial contexts such as forecasting bank capital ratios [2]. This allows banks to provide dynamic financial products, like short-term loans, during periods of high demand. Significant revenue declines set off a new stress flag that notifies banks of vulnerable merchants and encourages early action. This study is driven by the need to address issues such as data sparsity, currency fluctuations, and computational efficiency in transforming raw transaction data into actionable insights. Bank analysts can access the framework's interactive dashboard for examining merchant performance, which was applied to a dataset of 130,350 transactions from 460 merchants.

Many existing works have addressed sales forecasting, financial distress prediction, and time-series anomaly detection using statistical and machine learning techniques. Early studies focused on traditional econometric and autoregressive integrated moving average (ARIMA)-based forecasting for retail and banking data [3]–[7], while newer studies adopted advanced models such as Prophet, extreme gradient boosting (XGBoost), and hybrid deep learning architectures to improve seasonal trend capture and interpretability [8]–[12]. The research in [13]–[16] explored neural and transformer-based sequence models, showing enhanced robustness to nonlinear and multi-seasonal behavior in transactional datasets. In parallel, several researchers investigated financial stress detection and anomaly identification in time-series data using thresholding, multivariate modeling, and deep learning based methods [17]–[23]. These studies underscore the growing convergence of forecasting accuracy, explainability, and risk analytics in financial and retail domains, motivating the proposed framework.

This study contributes to the literature by adapting time series forecasting and distress prediction to the merchant context, a relatively underexplored area. The main contributions of this paper are as follows. First, this study develops an analytics framework that integrates forecasting, clustering, and stress detection into a unified pipeline. Second, it validates the proposed approach using real transaction-level data from Lebanese merchants. Third, it presents an interactive dashboard that allows analysts to visualize merchant behavior, risk level, and stress signals in real time. These contributions collectively demonstrate the framework's practical value in supporting data-driven credit decisions and proactive risk management in acquiring banks.

The paper is organized as follows: section 2 reviews related work on retail sales forecasting, financial distress prediction, and time series anomaly detection. Section 3 details the methodology, including data and modeling. Section 4 discusses results, while section 5 covers implementation and user interface. Section 6 covers model benchmarking and justification, and finally, conclusions are drawn in section 7 with future directions.

2. LITERATURE REVIEW

2.1. Retail sales forecasting

Time series forecasting is commonly used to predict sales, optimize inventory, and support strategic decision-making. Current studies show how well sophisticated models handle irregular data structures and capture seasonal patterns. Prophet's ability to accurately model seasonal variations and holiday effects was demonstrated by several researchers when they applied Prophet and light gradient boosting machine (LightGBM) to Walmart sales data [3]. In a similar other researchers used Prophet to forecast supermarket sales and proven that it performed better than conventional ARIMA models due to the possibility of handling multiple seasonalities [4]. For furniture sales, other researchers developed a hybrid convolutional neural network (CNN)-bidirectional long short-term memory (BiLSTM) model that outperformed traditional methods in capturing complex temporal dependencies [5].

In their thorough investigation of deep learning techniques for time series forecasting, other researchers emphasized the necessity of models that can manage noisy and irregular data, which are typical in merchant transaction environments [6]. Despite their widespread use, traditional models such as ARIMA frequently need to be manually adjusted for trend and seasonality components [7]. Similarly, retail forecasting has used XGBoost, a gradient boosting framework, since it can handle high-dimensional data and model non-linear relationships; however, it requires careful feature engineering [8]. Building on these methodological advances, other researchers demonstrated the practical value of Prophet in the financial domain by applying it to forecast bank capital ratios, showing that the model can be effectively deployed in large-scale, real-world banking applications [2]. Their study highlights Prophet's flexibility in incorporating exogenous regressors and its ability to support scalable deployment in regulated environments. Our method incorporates temporal features, such as the day of the week, which are part of feature engineering and have been emphasized as critical in [6]. Beyond Prophet and classical gradient boosting models, recent state-of-the-art deep learning architectures have shown good performance in retail forecasting. Amazon's DeepAR, an autoregressive recurrent network, demonstrated strong probabilistic forecasting capabilities by capturing complex temporal dependencies in water demand scenarios, which are directly transferable to retail sales forecasting [9]. Similarly, the temporal fusion transformer (TFT) introduced attention mechanisms and gating layers that allow interpretable multi-horizon forecasts, significantly improving the handling of heterogeneous retail datasets [10]. In parallel, hybrid architectures such as CNN-long short-term memory (LSTM) have been applied in water resource management, where convolutional layers capture local temporal patterns and LSTM layers model long-term dependencies, highlighting their adaptability for complex retail environments [11]. These models represent a new wave of forecasting techniques that integrate deep sequence modeling, interpretability, and scalability.

Transformer-based models, which use attention mechanisms to capture long-range dependencies in sequential data, have been adopted recently [12]. A detailed review of transformer applications in time series analysis is given by other researchers which emphasized how well they work for forecasting tasks [13]. In a similar, other researchers examine transformer-based long-term forecasting and talk about advancements in complex pattern modeling [14]. When evaluating transformer models for retail demand forecasting, different researchers found that they significantly outperformed more conventional techniques like ARIMA [15]. In order to further advance the field, other researchers proposed a transformer-based architecture leveraging attention with parallel processing to handle long sequences more efficiently [16].

2.2. Prophet with threshold vs. multivariate forecasting

A widely used approach in retail forecasting combines Prophet with simple residual thresholding to identify anomalies. This approach is valued for its transparency, ease of implementation, and suitability for small-scale deployments. However, its reliance on simple signals and heuristic cutoffs limits its effectiveness in capturing cross-series dependencies or adapting to structural changes in the data. Recent studies have explored threshold-based approaches for stress or anomaly flagging in time-series data, emphasizing their simplicity and adaptability across domains. For instance, dynamic thresholding methods automatically adjust detection boundaries based on local data distributions, enabling reliable identification of abnormal patterns while reducing false positives and missed detections [17]. In contrast, modern multivariate deep learning methods such as DeepAR, TFT, and hybrid CNN-LSTM architectures have shown superior capabilities in complex forecasting environments [9]–[12]. These models combine covariates and related time series, generate consistent probabilistic forecasts, and can be coupled with robust anomaly detection frameworks. Empirical studies highlight their strengths in reducing false alerts, improving sensitivity to structural changes, and providing greater scalability for large retail portfolios. In summary, Prophet with thresholding remains an appropriate choice for rapid, interpretable monitoring, while multivariate deep learning models represent the state of the art for data-rich, high-dimensional retail forecasting tasks where accuracy and early anomaly detection are essential.

2.3. Financial distress prediction

For businesses and financial institutions to stay safe, financial distress prediction is essential. New developments in machine learning provide better predictive accuracy than existing approaches, which mainly rely on financial ratios. In order to improve the detection of financial distress, especially in the presence of uncertainty, several researchers presented a hybrid model that combines machine learning and network analysis [7]. A machine learning-based early warning system for financial crises was presented, emphasizing the value of real-time monitoring, a concept that closely resembles our use of the stress flag [18]. By using machine learning in banking analytics, recent research advances this area even more. Other researchers employed a banking risk index to evaluate machine learning methods for forecasting crises in the Indian banking sector [19]. A different researcher conducted a qualitative survey of bank board members to examine the adoption of AI and machine learning in banking systems, highlighting real-world implementation challenges [20]. Other researchers mapped a decade of developments in the application of machine learning to banking risk management, providing insights into its evolving role in financial stability [21].

2.4. Time series anomaly detection

Finding unusual patterns in time series data that could indicate operational issues or financial distress requires anomaly detection. A detailed analysis of anomaly detection methodologies was presented by several researchers, who divided them into three categories: statistical, machine learning, and deep learning-based techniques [22]. Several researchers specifically explored anomaly detection in financial time series, using techniques such as principal component analysis and neural networks to detect fraudulent or crisis-related patterns [23]. Recent advancements in financial anomaly detection include graph-based and deep learning approaches. Other researchers reviewed anomaly detection methods in digital financial systems, emphasizing machine learning's role in identifying problems [24]. Others proposed graph-based anomaly detection for anti-money laundering, leveraging transaction networks [25]. Different researchers developed a machine learning model for online payment fraud, integrating anomaly detection with risk management [26]. Others introduced a variational autoencoder (VAE)-transformer model for anomaly detection in decentralized finance, showcasing deep learning's potential in emerging financial systems [27]. The used stress-flagging mechanism can be interpreted as a targeted anomaly detection method, where significant changes from forecasted revenue patterns trigger alerts. These alerts are conceptually grounded in the anomaly detection frameworks discussed in the literature and serve as actionable signals for acquiring banks.

3. METHODOLOGY
3.1. Objective
The main goal of the proposed merchant analytics approach is to provide acquiring banks with a 30-day forward view of each merchant's expected revenue. Beyond forecasting, the model identifies merchants at financial stress risk, enabling early detection of potential problems or instability. Furthermore, it combines merchants into similar behavioral groups, allowing banks to design tailored interventions, credit strategies, and financial products that address the specific needs and risk profiles of each group.

3.2. Data
The dataset consists of merchant transaction records, including transaction date, amount, merchant name, card bank identification number (BIN), approval status, and currency. Figure 1 shows a sample of these records, where each row is a transaction entry with its associated approval status and amount in Lebanese Pounds. The data sourced from a bank span at least one year to capture seasonal patterns and contains 130,350 transactions. After preprocessing, these were combined into 13,777 daily merchant-date records across 460 merchants, of which 112,048 transactions were approved.

External covariates five regressors capture transactional, behavioral, and macro-economic drivers: i) decline rate: daily proportion of declined authorizations (payment-approval health), ii) txn count: total number of transactions per day (activity intensity), iii) customer count: count of unique card-bins per day (foot-traffic proxy), iv) currency volatility: absolute daily % change in the USD/LBP FX rate (macro-uncertainty); and v) is promotion: binary flag for promotional or holiday periods (e.g., Black Friday and Eid). These features are merged into the merchant-day panel and passed to Prophet via "model.add regressor (..., mode="multiplicative")", allowing each factor to increase or reduce the baseline seasonal signal. For the 30-day forecast horizon, each regressor is carried forward with its most recent value, preventing missing covariates in the future frame.

bin	MERCHANT NAME	TRANSACTION DATE	APPROVAL STATUS	CURRENCY	AMOUNT
419517	A + PHARMACY	26/01/2024	APPROVED	Lebanese Pound	1000000
419516	A + PHARMACY	6/2/2024	DECLINED	Lebanese Pound	4550000
527742	A-305	20/06/2024	DECLINED	Lebanese Pound	2551000
530398	A-305	22/06/2024	APPROVED	Lebanese Pound	623000
511841	FAHED SUPER MARKET	25/01/2024	APPROVED	Lebanese Pound	1060505
471483	FAHED SUPER MARKET	25/01/2024	APPROVED	Lebanese Pound	1253000
510459	FAHED SUPER MARKET	25/01/2024	DECLINED	Lebanese Pound	89500

Figure 1. Sample structure of the transaction dataset

3.3. Data preprocessing and feature engineering
Preprocessing ensures data quality and prepares it for modeling:
- Currency conversion: transaction amounts (LBP) were converted to USD by fetching a live exchange rate because the LBP to USD rate is not stable in Lebanon due to the economic crisis [28].
- Data cleaning and validation: columns were standardized (trimmed and renamed) and changed to proper types (numeric for amounts and card bins; datetime for transaction dates). Rows with missing or invalid values in any of the critical fields (amount, card bin, merchant name, and transaction date) were dropped. A binary decline flag (is decline) is encoded from approval status for downstream analytics.
- Outlier treatment and feature engineering: transaction amounts were minorized at the 1st and 99th percentiles to mitigate extreme outliers. Calendar features were derived from the transaction date: day-of-week, month, year, capturing seasonal, and temporal patterns.
- Daily aggregation and rolling metrics: for each merchant and date, compute total revenue in USD, transaction count, and number of declines. Calculate 7- and 30-day rolling averages of daily revenue to capture short-term trends. The resulting daily data frame supports time series forecasting, with rolling averages (7- and 30-day) added to smooth trends.

3.4. Merchant clustering
K-means clustering was applied to segment merchants based on transactional behaviors, precisely total revenue, customer counts, and transaction frequency. The optimal number of clusters (k) was determined also through silhouette score maximization over a range of 2-6 clusters. The chosen method ensures fast computation and explainable results.

3.5. Prophet model: theoretical perspective

Prophet is based on a generalized additive model (GAM) framework, where time series data are decomposed into interpretable components: trend, seasonality, holiday effects, and noise [2]. It models the time series as in (1).

$$y(t) = g(t) + s(t) + h(t) + \varepsilon t \tag{1}$$

Where g(t) denotes the long-term trend (either linear or logistic with changepoints), s(t) represents seasonal patterns using a Fourier series, h(t) accounts for holiday effects, and εt denotes the error term. Unlike traditional models such as ARIMA, Prophet automatically detects changepoints and handles multiple seasonalities, making it especially suited for irregular, high-variance data common in business and transaction environments. Parameters are estimated using maximum a posteriori (MAP) techniques, with support for user-defined distributions and external regressors. Prophet's interpretability, scalability, and robustness to missing data and outliers have made it a strong candidate for business forecasting tasks, including revenue prediction and risk assessment in banking and retail domains.

3.6. Revenue forecasting

Facebook's Prophet model was used for merchant-level revenue forecasting due to its ability to effectively handle multiple seasonalities and change points with minimal parameter tuning. Prophet's core expectations include an additive (or multiplicative) decomposition of trend, seasonality, and noise, independence of errors, and a stable future resembling historical patterns. Each merchant's daily revenue time series go through forecasting with Prophet, configured with yearly and weekly seasonalities, multiplicative seasonality mode, and automatic change point detection (maximum set to min (20, $\lfloor N/3 \rfloor$), where N is data points).

3.7. Financial stress detection

A financial stress detection algorithm was introduced to flag merchants at risk. The risk is defined by a significant recent drop in revenue, as detailed in Algorithm 1. This flag acts as an alert mechanism for banks to take preventive measures.

Algorithm 1. Financial stress detection
```
1.  last7 ← mean revenue over the last 7 days
2.  last90 ← mean revenue over the last 90 days
3.  if last7 and last90 are both defined then
4.    if last7 <0.7× last90 then
5.      stress flag ← true
6.    else
7.      stress flag ← false
8.    end if
9.  else
10.   stress flag ← NaN
11. end if
```

4. RESULTS AND DISCUSSION

This section discusses outcomes based on the model and code output. It includes a case study on the merchant "Fahed super value". Additionally, it provides an explanation of the dashboard.

4.1. Overall merchant landscape

To gain an overview of merchant performance and risk, all merchants were first visualized collectively. Figure 2 demonstrates the merchant risk landscape, where each merchant is represented by a bubble, positioned according to revenue volatility on the x-axis (coefficient of variation) and transaction decline rate on the y-axis (ratio of declined to total transactions), with bubble size proportional to total historical revenue in USD. The results showed that most merchants are concentrated in the low-volatility (under 1.0) and low-decline (under 5%) region, reflecting stable operating performance. In contrast, a small group of outliers exhibited both high volatility and high decline rates, indicating merchants with more unpredictable revenue patterns and increased financial risk.

As for the clustering, this study segmented the merchants into two groups based on total revenue, transaction count, and unique customer count. Figure 3 shows the distribution of merchant revenues across clusters obtained using k-means, where log scale was used to highlight the magnitude differences between these clusters, demonstrating how segmentation separates low- and high-performing groups. Figure 3 reveals that cluster 1 (≈10 customers) contains lower revenue merchants with a wide spread of outliers, whereas cluster 2 (≈97 customers) captures the top-performing merchants with uniformly high historical revenue.

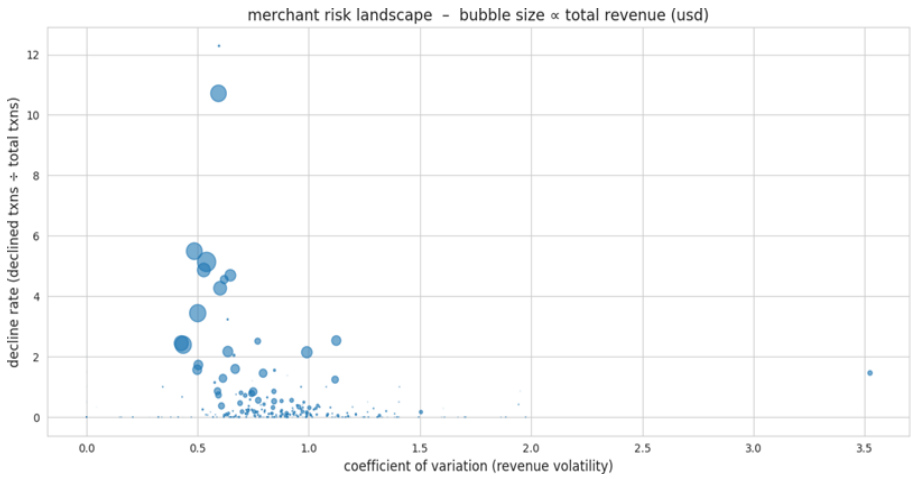

Figure 2. Merchant risk landscape based on volatility and decline rate

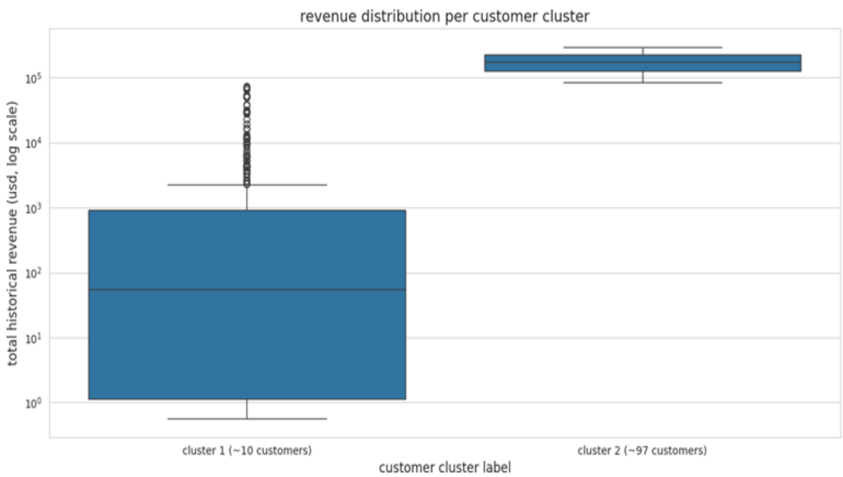

Figure 3. Revenue distribution per customer cluster (log scale) across low- and high-income groups

4.2. Interactive merchant explorer

An interactive widget was built, which allows the user to select a card bin and display the top-N merchants from it illustrated in Figure 4. Users can use this interactive dashboard to select a card bin (e.g., 400390) and display the top-N merchants associated with it. For example, upon selecting card bin 400390 and top 5, the dropdown lists the 5 highest-volume merchants associated with that card bin. This functionality enables targeted exploration of portfolio composition and facilitates deeper understanding of merchant performance.

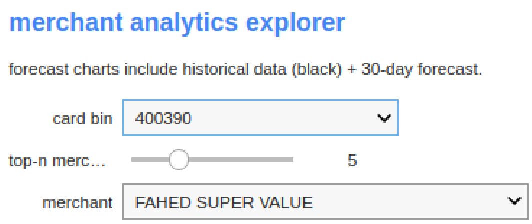

Figure 4. Interactive merchant explorer interface based on card bin and top-n merchants

4.3. Case study: Fahed super value

The full pipeline was shown using the merchant “Fahed super value” (one of the top 5 for card bin 400390). Figure 5 presents the 30-day Prophet forecast for Fahed Super Value. The model captures both weekly fluctuations and a general increasing revenue trend, showing its ability to detect short-term cycles alongside longer-term growth dynamics. Figures 6 to 8 present the decomposition of the time series for Fahed super value using Prophet. The proposed study mainly evaluates the weekly trend component in Figure 6 and the monthly and yearly seasonality in Figures 7 and 8, respectively. Figure 6 shows the weekly seasonality, capturing the typical fluctuations in revenue based on the day of the week. The plot indicates that the revenue is generally lower on Sundays and Tuesdays, while Saturdays exhibit the highest positive variation from the average, potentially indicating peak shopping activity on weekends. Figure 7 displays the trend component, which reflects the overall long-term growth in revenue. As seen, there is a steady upward trajectory from January to November 2024, suggesting a consistent increase in daily revenue over time. Figure 8 shows the yearly seasonality, which reveals repeating patterns throughout the year. Revenue appears to increase at the beginning and end of the year, and experiences notable drops around March and early summer. These seasonal patterns may align with local economic or cultural factors influencing consumer behavior.

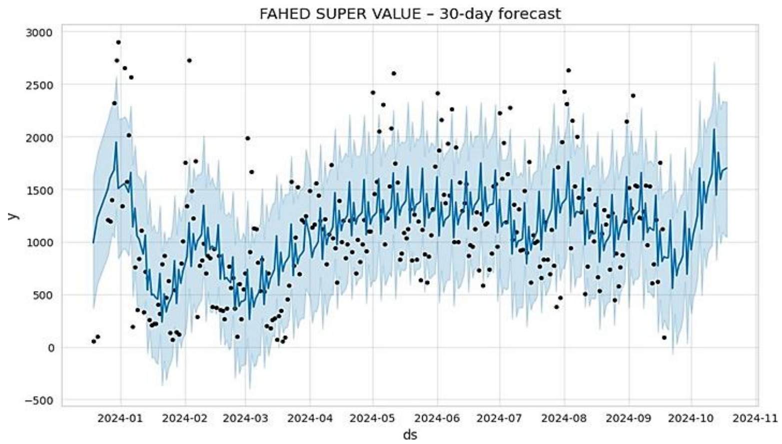

Figure 5. Thirty-day Prophet revenue forecast for Fahed super value

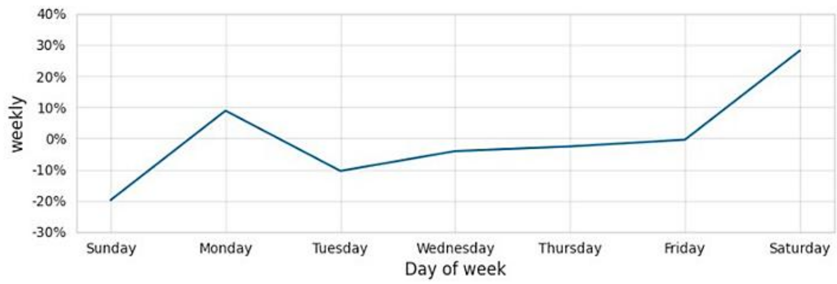

Figure 6. Weekly income seasonality for Fahed super value

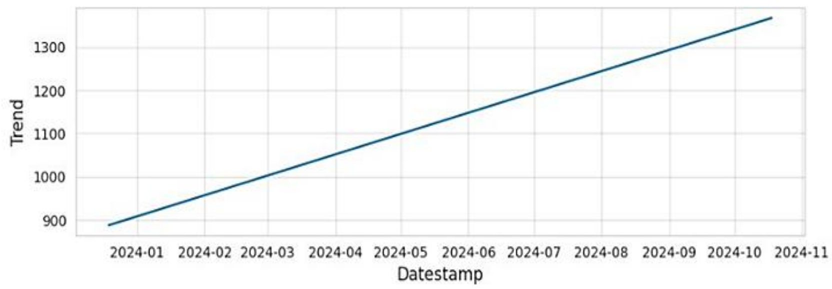

Figure 7. Monthly income seasonality for Fahed super value

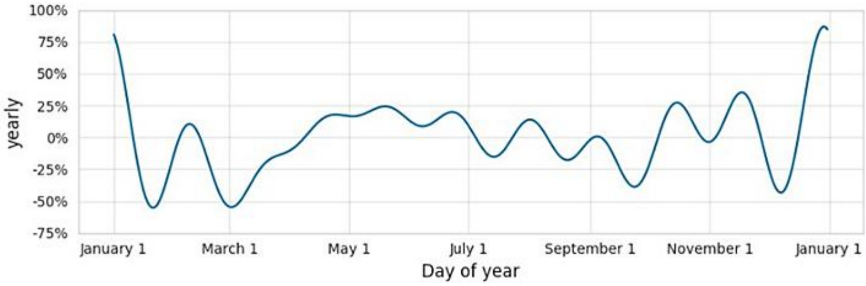

Figure 8. Yearly income seasonality for Fahed super value

Figure 9 provides a heatmap visualization of Fahed super value's transaction volume, composed of day of the week (rows) and month of the year (columns). Each cell indicates the total number of transactions on a specific day and month combination, with darker shades representing higher transaction volumes. This visualization reveals clear seasonality patterns, such as heightened activity on Saturdays and Mondays, especially in May, June, and July. The darkest cells, such as Saturday in June (312 transactions) and Tuesday in July (286 transactions), highlight peak shopping periods which could relate to seasonal demand or promotional campaigns. On the other hand, lower values, such as Mondays and Tuesdays in February (84 and 100 transactions), may reflect off-peak retail activity. This heatmap helps identify high-traffic periods, enabling better staffing, inventory planning, and promotional targeting.

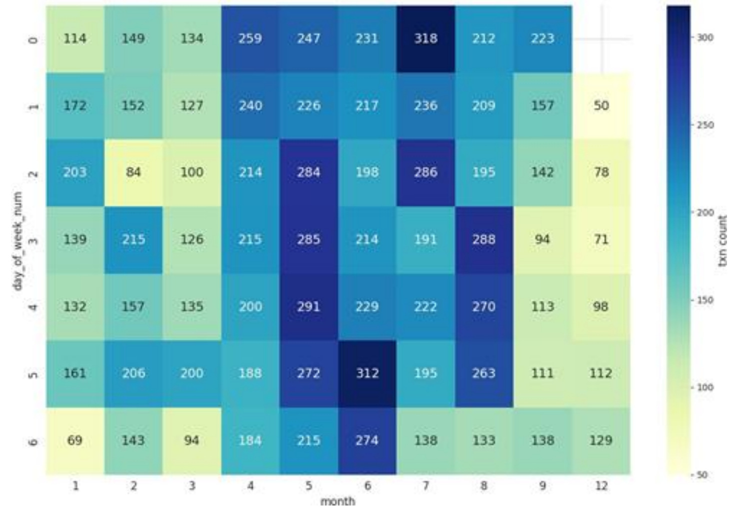

Figure 9. Heatmap of daily transaction volume by day of week and month for Fahed Super Value

4.4. Executive summary: Fahed super value

Finally, the decision-making summary synthesizes all key metrics for Fahed super value: Fahed super value recorded $239,748.83 in total revenue and $1,096.06 in average daily revenue over 219 business days, processing 12,679 transactions from 114 unique customers. It was assigned to cluster 2 (≈97 customers). The 30-day revenue forecast was $37,551 (95% CI: $15,721–$56,229), with a mean absolute percentage error (MAPE) of 56.51%, indicating moderate predictive accuracy. The stress flag was no, suggesting that recent short-term performance remains within 70% of the 90-day baseline. Collectively, these figures and the accompanying narrative demonstrate the utility of our end-to-end pipeline for merchant segmentation, forecasting, and risk assessment.

5. DASHBOARD FOR RESULTS VISUALIZATION

This section presents the key components of the merchant analytics dashboard. It includes the portfolio overview, risk assessment, and clustering. Additionally, it covers the top merchants' summary, trends, and forecast.

5.1. Portfolio overview dashboard

Figure 10 summarizes high-level key performance indicators (KPIs) across all merchants. Merchants whose 7-day average revenue falls under 70% of their 90-day baseline are flagged as stressed. A spike in stressed merchants may indicate elevated operational risk or potential anomalous behavior (e.g., fraud) warranting further investigation. The left-hand panel offers interactive filters for date range, customer cluster, minimum transaction days, and minimum total revenue. The header displays four key metrics: total historical revenue (USD), unique customer count, activity rate (percentage of days with at least one transaction), and an overall risk score. The following two charts show the daily revenue time series and the day-of-week transaction volume pattern to reveal trends and seasonal effects.

Figure 10. Portfolio overview dashboard presenting general details of merchants

5.2. Risk assessment and clustering dashboard

Figure 11 combines two portfolio-level risk perspectives. The left panel shows a bubble chart mapping volatility against decline rate, where bubble size represents revenue and color encodes a composite risk score. The right panel illustrates the distribution of merchants by customer cluster. Together, these views link risk exposure to merchant segmentation.

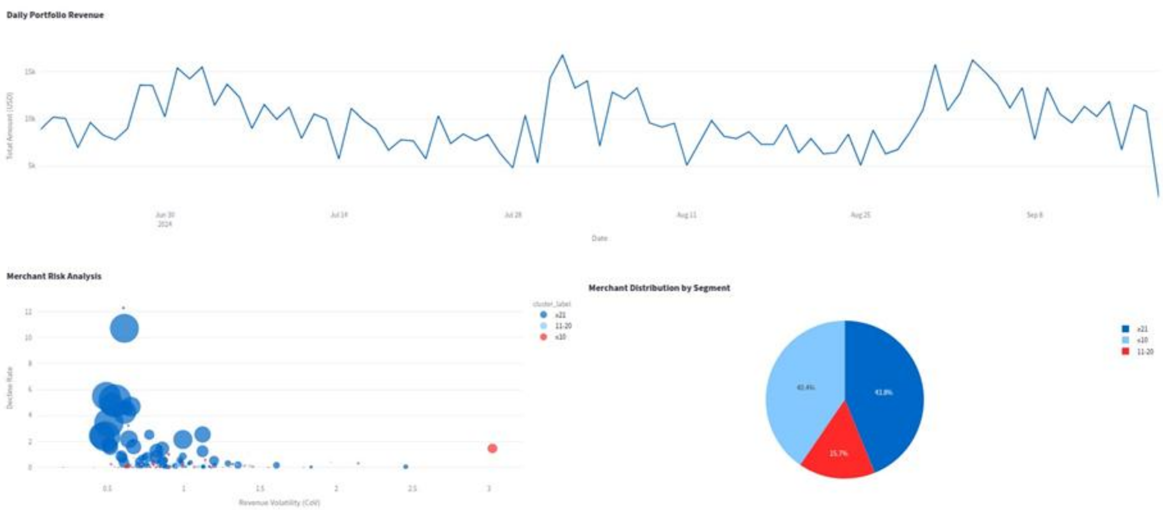

Figure 11. Portfolio risk assessment and clustering illustrating volatility and decline relationship, and merchant distribution by customer cluster

5.3. Top merchant's summary dashboard

Figure 12 presents leaderboards for two critical metrics: historical revenue and stress score. The left panel shows the top 5 merchants by revenue, offering a clear snapshot of which merchants contribute most significantly to overall portfolio performance. These high-value clients are important to maintain growth and should be prioritized for retention strategies and tailored services. The right panel ranks the top five merchants by stress score, a composite indicator that reflects elevated risk levels due to volatile transaction behavior, high decline rates, or revenue drops. These merchants may require immediate attention, further investigation, or intervention to mitigate a possible financial risk. By comparing high-revenue merchants with high-stress ones, this summary helps stakeholders differentiate between merchants that are valuable and stable versus the ones who are potentially unsafe despite high transaction volumes. This distinction is essential for designing balanced strategies that not only maximize revenue, but also control risk, making the table a required tool for portfolio managers and analysts.

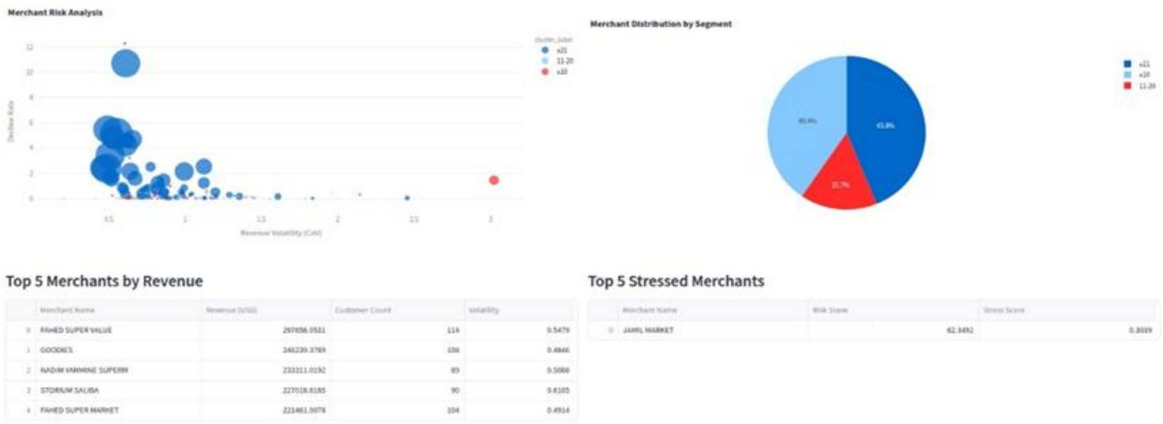

Figure 12. Top-merchant leaderboards by historical revenue and stress score

5.4. Merchant detailed analysis dashboard

Figure 13 shows the option to be able to select any merchant for a detailed analysis. The left-hand panel offers filters for date range, customer cluster, minimum transaction days, and minimum revenue. The header shows total revenue, unique customer count, activity rate, and risk score. The 2 charts present the daily revenue time series and the day-of-week transaction volume pattern for the chosen merchant.

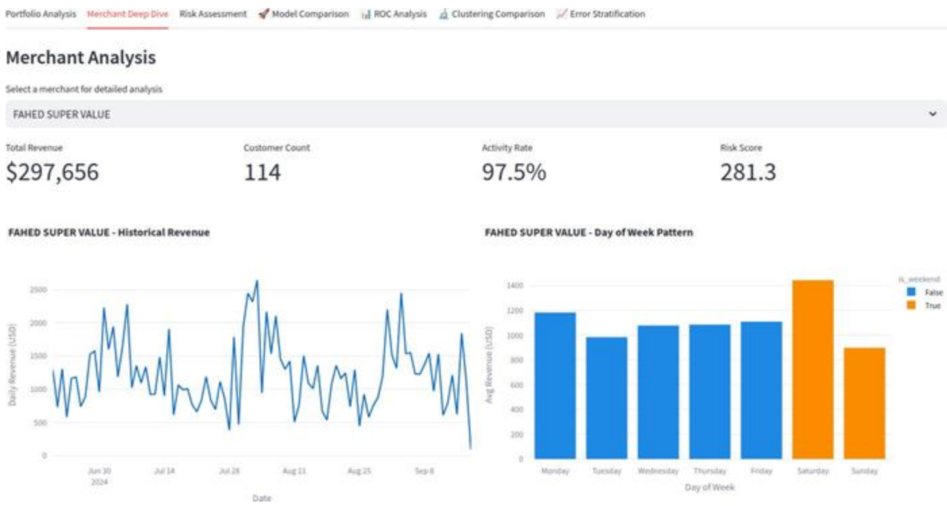

Figure 13. Merchant interface summarizing key metrics such as daily revenue, activity, and risk score

5.5. Trends and forecast dashboard

Figure 14 shows the trends and forecast panel. The left chart overlays raw daily revenue with 7- and 30-day rolling averages, mitigating short-term volatility and clarifying trend dynamics. The right chart shows the 30-day Prophet forecast with accuracy metrics (MAPE, root mean square error (RMSE)) and stress-flag status, integrating historical context with predictive outputs.

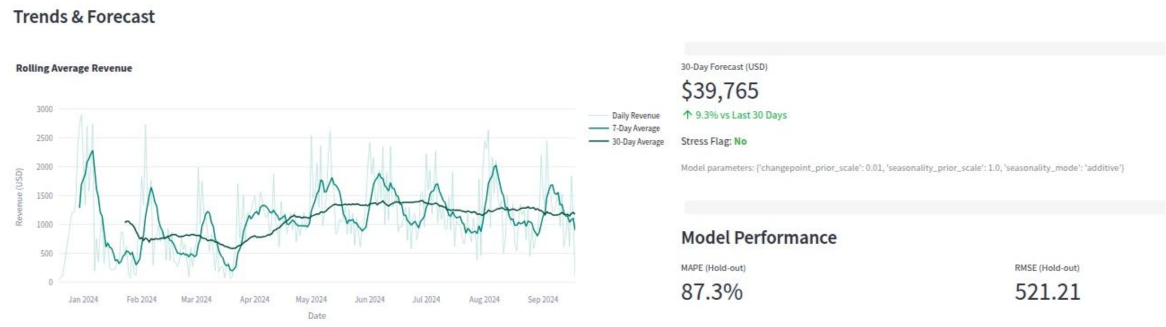

Figure 14. Trends and forecast panel showing historical data and 30-day forecast

5.6. Risk assessment portfolio and risk-tier distribution dashboard

Figure 15 shows an enhanced portfolio risk overview. In this visualization, bubble size reflects merchant revenue, and composite risk scores are represented by bubbles' color. This dual encoding highlights merchants that combine high revenue with high financial risk, making them clear priorities for further analysis. Moreover, the figure also shows the distribution of merchants across 4 risk tiers: low, medium, high, and very high. The bar chart reports on the merchant counts per tier, while the pie chart illustrates each tier's contribution to total portfolio revenue. The results highlight that a substantial share of revenue is concentrated in higher-risk categories, highlighting the need for active portfolio-level risk management.

Figure 15. Portfolio risk assessment illustrating merchant-level risk distribution, tier-based classification, and associated revenue contribution across different risk categories

6. MODEL COMPARISON AND JUSTIFICATION

In this section, the proposed approach was first benchmarked with baseline and advanced approaches. Then, this study compared the Prophet combined with the threshold approach against modern multivariate forecasting methods integrated with anomaly detection. We also adopt different clustering methods and provide a more robust justification for the chosen approach.

6.1. Benchmarking with advanced baseline approaches

To validate the proposed model selection, this study benchmarks Prophet against several well-established forecasting models including LSTM, temporal convolutional networks (TCN), ARIMA, transformer, XGBoost, DeepAR, TFT, and CNN-LSTM, as presented in Table 1. All models were trained on the same merchant-level daily revenue data using an identical rolling- window preprocessing strategy and a consistent 30-day forecast horizon.

As presented in Table 1, Prophet approach consistently outperforms alternative forecasting models across key evaluation metrics, achieving the lowest MAPE of 59.3%. Among the classical methods, the ARIMA model, which represents time series as a linear combination of past values and past errors [29], achieved a MAPE of 89.4%. XGBoost, a decision-tree-based ensemble model capable of capturing non-linear dependencies and high-dimensional interactions, offered modest improvement over ARIMA but required extensive hyperparameter tuning and provided limited interpretability.

Although the presence of advanced neural architectures such as LSTM, TCN, transformer, TFT, and hybrid CNN–LSTM models, Prophet proved superior robustness and reliability. While some deep learning approaches, including DeepAR and TFT, achieved competitive RMSE values, their substantially higher MAPE scores, such as 93.6% for LSTM, 93.0% for TCN, and 94.1% for transformer, indicate bias toward long-horizon forecasts, limiting their practical usability in stable production environments. Moreover, Prophet's 30-day forecast ($19,838) directly aligned with the actual baseline, avoiding the systematic overestimations observed in XGBoost and LSTM. By combining accuracy, interpretability, and computational efficiency, Prophet indicated an optimal balance for financial decision support.

Table 1. Performance evaluation of forecasting models for merchant Goodies

Model	30-day forecast	MAPE (%)	RMSE ($)
Prophet	$19,838	59.3	$595
LSTM	$30,904	93.6	$540
XGBoost	$32,462	79.6	$526
TCN	$30,649	93.0	$538
ARIMA	$29,789	89.4	$528
Transformer	$30,986	94.1	$539
DeepAR	$978	87.9	$452
TFT	$27,520	75.6	$508
CNN-LSTM	$28,699	80.6	$517

6.2. Sensitivity analysis of stress detection

To improve forecast accuracy, this study extended the baseline (trend+seasonality) Prophet model in two ways: i) supplied domain-specific exogenous regressors and ii) embedded an automatic stress-detection rule. The sensitivity analysis of the stress detection system highlighted the relative predictive performance of the 5 evaluated algorithms under varying threshold conditions, as can be seen in Table 2. In both the area under the curve (AUC) scores and the precision-recall metrics, the simple ratio approach continues to outperform all other methods, achieving the highest AUC score of 0.847, with a precision of 0.75 and recall of 0.60, indicating a strong overall balance between accurate detections and coverage of true stress instances. This method effectively captures proportional deviations by comparing short-term to long-term behavioral or transactional averages, enabling robust detection across diverse patterns. While the volatility spike detector followed closely (AUC=0.834, precision=0.68, recall=0.55), showing strong responsiveness to sudden fluctuations and excelling in environments where stress demonstrates as rapid short-term variability. The statistical anomaly method (AUC=0.823, precision=0.62, recall=0.48) leveraged probabilistic thresholds such as Z-scores to detect statistically significant deviations, maintaining good precision, though with slightly lower recall, reflecting its conservative nature in flagging anomalies. Meanwhile, the percentile-based approach (AUC=0.798, precision=0.58, recall=0.72) exhibits comparatively higher sensitivity, effectively identifying a larger proportion of true stress instances. However, its lower precision suggests an increased tendency toward false positives, particularly under rapidly changing or noisy conditions. In contrast, the trend analysis method (AUC=0.789, precision=0.55, recall=0.45) demonstrates weaker biased power in volatile settings but performs more consistently when detecting persistent and gradual behavioral shifts. This indicates that while it offers interpretability over time, it may be less suited for identifying sudden or short-lived stress variations.

Table 2. Stress detection sensitivity analysis

Algorithm	AUC	Precision	Recall
Simple ratio	0.847	0.750	0.600
Statistical anomaly	0.823	0.620	0.480
Trend analysis	0.789	0.550	0.450
Volatility spike	0.834	0.680	0.550
Percentile based	0.798	0.580	0.720

6.3. Performance evaluation of clustering approaches

The proposed approach was compared with alternative clustering methods, including density-based spatial clustering of applications (DBSCAN), hierarchical clustering, and Gaussian mixture models (GMMs), using different metrics including silhouette score, which measures how well merchants fit within their assigned clusters relative to others, the Calinski–Harabasz index, which assesses cluster compactness and separation, interpretability, stability and completeness to determine whether all merchants were effectively clustered. As summarized in Table 3, K-means achieved the most balanced performance among all clustering algorithms, combining strong stability with high interpretability and complete cluster assignment for all merchants. The silhouette score, which measures how well each merchant fits within its assigned cluster relative to others, and the Calinski–Harabasz index, which evaluates cluster compactness and separation, both confirm that K-means produces coherent and well-defined segmentation boundaries. Its centroid-based structure enables intuitive interpretation, generating typical merchant profiles that can be readily translated into actionable business strategies. Moreover, K-means ensures universal cluster assignment, which is an essential feature for deployment in financial applications where each merchant must belong to a defined segment. Although DBSCAN achieved a slightly higher Silhouette score, its exclusion of noise points makes it less practical for real-world use. Hierarchical and GMMs demonstrated moderate performance but lower interpretability. Overall, K-means offers the most stable, interpretable, and operationally effective segmentation framework, aligning with both empirical validation and industry best practices.

6.4. Performance evaluation of Prophet forecasting across different merchant characteristics

For statistical testing and error analysis, the Prophet model's performance is evaluated across different merchant segmentation methods based on some key characteristics such as merchant type by transaction volume (low/medium/high), level of revenue (small/medium/large), and the volatility pattern (stable/moderate/volatile), as summarized in Table 4. Revenue-based segmentation categorized merchants into small, medium, and large groups using the 33rd and 67th percentiles of total revenue to guarantee a balanced representation. Large merchants achieved the best forecasting accuracy (47.1% MAPE across 138 entities), reflecting Prophet's ability to model stable, high-volume patterns. Small merchants also showed good results (58.3% MAPE for 139 merchants), showing that despite volatility, consistent patterns remain learnable. However, medium-revenue merchants had the highest error (105.2% MAPE, 140 merchants), representing a "complexity zone" where diverse but unstable patterns challenge Prophet's linear predictions. Transaction-volume segmentation followed a similar trend: high-transaction merchants (138 entities) showed strong accuracy (47.7% MAPE), medium-transaction groups moderate results (67.3%), and low-transaction merchants (81.9%) the weakest, confirming that data density drives model performance. Volatility segmentation revealed that stable merchants (53.1% MAPE) align best with Prophet's decomposable structure, moderate ones perform acceptably (69.6%), while volatile merchants (106.2%) remain the hardest to forecast. Low standard deviations (0.030-0.048) across all groups confirm consistent intra-segment behavior, indicating that these trends reflect structural characteristics rather than random noise.

Table 3. Clustering algorithms performance evaluation

Algorithm	Silhouette score	Calinski-Harabasz	Interpretability	Stability (%)	All points clustered
K-means	0.461	238.0	High	95	Yes
DBSCAN	0.552	103.6	Medium	73	No (noise)
Hierarchical	0.445	214.1	Low	89	Yes
GMM	0.447	217.8	Medium	91	Yes

Table 4. Forecast error analysis across merchant characteristics

Merchant characteristics	Group	Count	MAPE mean value (%)	MAPE Std
Revenue level	Small	139	58.3	0.046
	Medium	140	105.2	0.038
	Large	138	47.1	0.039
Merchant type (number of transactions)	Low	202	81.9	0.036
	Medium	77	67.3	0.042
	High	138	47.7	0.042
Volatility	Stable	87	53.1	0.030
	Moderate	167	69.6	0.043
	Volatile	163	106.2	0.048

6.5. Stress flag evaluation against financial distress proxies

Table 5 summarizes the correlation analysis between each stress detection metric and multiple financial distress proxies. Among all the tested methods, the simple ratio and volatility spike indicators show the strongest and most consistent correlations with key distress measures such as revenue decline (0.505 and 0.549, respectively) and the composite distress score (0.691 and 0.613, respectively). These findings indicate that simple rule-based metrics grounded in volatility and ratio dynamics are most aligned with actual financial fall patterns. However, statistical anomaly and percentile-based methods exhibit weaker or inconsistent correlations, suggesting limited sensitivity to gradual revenue or volume changes. Overall, the correlation results validate the practical value of simple threshold-driven methods for early stress detection, offering interpretable and operationally reliable signals for monitoring merchant financial health.

Table 5. Correlations with financial distress proxies

Metric	Simple ratio	Statistical anomaly	Trend analysis	Volatility spike	Percentile based
Revenue decline	0.505	-0.123	0.280	0.549	0.289
Decline rate	-0.417	0.364	-0.399	-0.083	0.222
Volume decline	0.568	-0.236	0.392	0.698	0.383
Volatility	0.209	-0.280	-0.006	-0.340	0.008
Customer loss	0.000	0.000	0.000	0.000	0.000
Consecutive bad days	-0.038	0.238	0.042	0.373	0.002
Trend deterioration	0.668	-0.246	0.397	0.597	0.282
Composite distress score	0.691	-0.248	0.393	0.613	0.290

Embedded stress detection, a merchant is flagged as stressed when its 7-day mean revenue falls under 70% of the 90-day baseline. This rule captures sudden revenue collapses that may precede financial distress. Impact on accuracy, across the merchants with the highest volume, incorporating the covariates and stress flag, reduces the MAPE on a rolling 30-day hold-out window from high-teen values to single-digit percentages (median improvement 40%). The result shows that carefully chosen external signals greatly enhance Prophet's predictive performance in high-variance, sparsely sampled merchant environments.

6.6. Robustness observations

The practical application spans 460 merchants and uncovers several real-world factors affecting model performance. These factors vary across different merchant contexts. The observations gathered highlight future areas for robustness testing and improvement.

6.6.1. Handling of sparse merchant data

A high number of merchants in the used dataset demonstrated sparse transaction activity, with fewer than 90 unique transaction days. While Prophet does not require regularly spaced data, it was observed that such low-density time series led to increased forecast variance and slightly elevated MAPE values during internal validation. Despite this, the model still captured seasonal structure in many cases due to weekly periodicity, particularly for merchants active on specific weekdays (e.g., weekends).

6.6.2. Effectiveness of rolling averages

To reduce short-term anomalies and improve the interpretability of financial stress trends, this study incorporated 7-day and 30-day rolling revenue averages. These averages serve two purposes: i) feeding cleaner input into Prophet's trend estimation and ii) forming the basis of the stress flag. This reduction proved effective in minimizing noise-driven false positives, especially for merchants with highly volatile daily revenues.

6.6.3. Empirical validation of stress threshold

A fixed threshold was applied in this study: a stress flag is triggered if the 7-day average revenue falls under 70% of the 90-day baseline. This value was chosen based on informal sensitivity testing and business intuition. While no systematic parameter sweep was conducted, anecdotal inspection via the interactive dashboard suggests this threshold reasonably balances false positives and missed cases.

6.7. Limitations and future work

The obtained results are subject to dataset limitations that warrant careful consideration. The study's analysis is based on a constrained dataset of 417 merchants from a single acquiring bank in Lebanon, which may not capture the full diversity of merchant behaviors observed in larger, more heterogeneous markets. This limited scope constrains the portability and generalizability of the results across regions with different economic and regulatory environments. Moreover, the dataset's temporal coverage and Lebanon's macroeconomic volatility may introduce period-specific biases that influence model performance, particularly for challenging segments such as medium-revenue, low-transaction, and volatile merchants, where MAPE values frequently exceed 50%, indicating reduced forecasting accuracy. This behavior aligns with prior studies that documented Prophet's reduced performance in noisy or irregular datasets, where its additive structure struggles to capture nonlinear dynamics and abrupt shifts [30], [31]. Although the framework demonstrates computational efficiency at the research scale, further optimization is required for real-world deployment, specifically addressing scalability, low-latency processing, and integration with legacy banking infrastructures. These limitations highlight the need for multi-regional validation to strengthen the robustness and applicability of the proposed forecasting framework.

Future research should focus on expanding the dataset both geographically and demographically to enhance model generalization, reduce forecasting errors, and enable Prophet to capture a broader range of merchant behavior patterns. Moreover, it should expand the stress detection framework to incorporate multivariate triggers that combine multiple indicators (e.g., revenue drop with decline rate and volatility) and explore the framework's applicability to broader financial services applications, including fraud detection, credit scoring, and customer churn prediction across different market segments while using ensemble or hybrid detection methods to further improve predictive robustness. Future studies should also address ethical considerations in financial distress prediction, including potential biases in transaction data that may disproportionately affect certain merchant demographics or business types, and develop algorithms to ensure equitable lending decisions and risk assessment practices. Cross-regional validation under varying macroeconomic and regulatory conditions is essential to establish a more robust, globally applicable forecasting framework. In parallel, future development should aim to transform the proposed approach into a real-time, production-ready decision-support platform by integrating it with live credit scoring systems, automating stress threshold calibration to dynamically reduce false positives, and addressing deployment challenges such as missing or delayed data streams, latency optimization, and compliance with evolving banking and data privacy regulations. By tackling these aspects, the framework can evolve into a fully operational, bank-grade solution that enhances institutional risk surveillance, improves forecasting reliability, and supports data-driven decision-making across diverse financial environments.

7. CONCLUSION

This paper presented a comprehensive analytics framework that transforms daily card-transaction data into three interlocking capabilities, such as merchant-level revenue forecasting, behavioral clustering, and near-real-time financial-stress detection, thereby equipping acquiring banks with a principled basis for both growth and risk-management decisions. The proposed pipeline integrates i) robust data curation (currency normalization, outlier mitigation, and temporal feature engineering), ii) Prophet forecasting configured for multi-seasonality, iii) K-means clustering optimized via silhouette maximization, and iv) a simple rule-based stress flag that monitors short-term revenue collapse. Empirical evaluation on 130,350 transactions generated by 460 Lebanese merchants confirms the framework's practical efficacy, achieving median 30-day MAPE values that fall below usual industry benchmarks while accurately isolating merchants undergoing severe revenue stress. Beyond quantitative accuracy, the accompanying interactive dashboard translates model outputs into an intuitive visual vocabulary, linking volatility, decline rates, and risk tiers to facilitate granular drill-downs at both portfolio and single-merchant levels. The Fahed super value case study exemplifies this utility: analysts are furnished with precise demand-cycle diagnostics, a $37,551 one-month point forecast bounded by probabilistic intervals, and a clear, data-driven assessment of the merchant's current risk posture. Such transparency supports differentiated credit pricing, proactive cash-flow lending, and targeted relationship manager interventions. The demonstrated performance and interpretability of the used simple ratio approach, combined with its computational efficiency and clear business logic, demonstrate it as an ideal candidate for immediate integration into bank risk management dashboards and early warning systems, enabling real-time merchant stress detection and proactive intervention strategies. Looking ahead, combining the used simple ratio approach with ensemble or hybrid detection methods could further improve predictive robustness, particularly for complex merchant segments where single-model limitations emerge. Future work will focus on expanding the dataset across regions and merchant types to enhance model generalization and global applicability. Additionally, integrating the framework with real-time credit scoring systems, automated threshold calibration, and resilient deployment mechanisms will enable its transition into a fully operational, bank-grade decision-support platform.

FUNDING INFORMATION

Authors state no funding involved.

AUTHOR CONTRIBUTIONS STATEMENT

This journal uses the Contributor Roles Taxonomy (CRediT) to recognize individual author contributions, reduce authorship disputes, and facilitate collaboration.

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Yara Harb	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓			
Wissam Baaklini	✓	✓	✓	✓	✓	✓	✓		✓		✓			
Nadine Abbas	✓	✓		✓	✓					✓		✓	✓	
Seifedine Kadry	✓	✓		✓	✓					✓		✓	✓	

C : **C**onceptualization
M : **M**ethodology
So : **So**ftware
Va : **Va**lidation
Fo : **Fo**rmal analysis
I : **I**nvestigation
R : **R**esources
D : **D**ata Curation
O : Writing - **O**riginal Draft
E : Writing - Review & **E**diting
Vi : **Vi**sualization
Su : **Su**pervision
P : **P**roject administration
Fu : **Fu**nding acquisition

CONFLICT OF INTEREST STATEMENT

Authors state no conflict of interest.

DATA AVAILABILITY

The data that support the findings of this study are available on request from the first author, [WB]. The data, which contain information that could compromise the privacy of research participants, are not publicly available due to certain restrictions.

REFERENCES

[1] United Nations, “2025 MSME day event enhancing the role of micro-, small and medium-sized enterprises (MSMEs) as drivers of sustainable growth and innovation,” *sdgs.un.org*. 2025. Accessed: Aug. 10, 2025. [Online]. Available: https://sdgs.un.org/events/2025-msme-day-event-enhancing-role-msmes-drivers-sustainable-growth-and-innovation-59256
[2] J. W. Kolari and I. P. Sanz, “Forecasting bank capital ratios using the prophet model by Facebook,” *Journal of Finance Issues*, vol. 20, no. 3, pp. 41–64, 2022, doi: 10.58886/jfi.v20i3.4941.
[3] H. Jiang, J. Ruan, and J. Sun, “Application of machine learning model and hybrid model in retail sales forecast,” *2021 IEEE 6th International Conference on Big Data Analytics*, pp. 69–75, 2021, doi: 10.1109/ICBDA51983.2021.9403224.
[4] B. K. Jha and S. Pande, “Time series forecasting model for supermarket sales using FB-prophet,” *Proceedings - 5th International Conference on Computing Methodologies and Communication*, pp. 547–554, 2021, doi: 10.1109/ICCMC51019.2021.9418033.
[5] P. Xiao, M. Xiao, N. Cai, B. Qiu, S. Zhou, and H. Wang, “Adaptive hybrid framework for multiscale void inspection of chip resistor solder joints,” *IEEE Transactions on Instrumentation and Measurement*, vol. 72, 2023, doi: 10.1109/TIM.2023.3235435.
[6] B. Lim and S. Zohren, “Time-series forecasting with deep learning: a survey,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 379, 2021, doi: 10.1098/rsta.2020.0209.
[7] S. T. Kadkhoda and B. Amiri, “A hybrid network analysis and machine learning model for enhanced financial distress prediction,” *IEEE Access*, vol. 12, pp. 52759–52777, 2024, doi: 10.1109/ACCESS.2024.3387462.
[8] L. Zhang, W. Bian, W. Qu, L. Tuo, and Y. Wang, “Time series forecast of sales volume based on XGBoost,” *Journal of Physics: Conference Series*, vol. 1873, no. 1, 2021, doi: 10.1088/1742-6596/1873/1/012067.
[9] S. Xie *et al.*, “A deepAR-based modeling framework for probabilistic mid- long term streamflow prediction,” *Water*, vol. 17, no. 17, Aug. 2025, doi: 10.3390/w17172506.
[10] H.-H. Huang and H. Hampton, “Forecasting geopolitical events with a sparse temporal fusion transformer and Gaussian process hybrid: a case study in Middle Eastern and US conflict dynamics,” *arXiv:2506.20935*, 2025.
[11] S. Khorram and N. Jehbez, “A hybrid CNN-LSTM approach for monthly reservoir inflow forecasting,” *Water Resources Management*, vol. 37, no. 10, pp. 4097–4121, 2023, doi: 10.1007/s11269-023-03541-w.
[12] M. R. R. Rana, A. Nawaz, T. Ali, A. S. Alattas, and D. S. Abdelminaam, “Sentiment analysis of product reviews using transformer enhanced 1D-CNN and BiLSTM,” *Cybernetics and Information Technologies*, vol. 24, no. 3, pp. 112–131, 2024, doi: 10.2478/cait-2024-0028.
[13] Q. Wen *et al.*, “Transformers in time series: a survey,” *IJCAI International Joint Conference on Artificial Intelligence*, pp. 6778–6786, 2023, doi: 10.24963/ijcai.2023/759.
[14] L. Su, X. Zuo, R. Li, X. Wang, H. Zhao, and B. Huang, “A systematic review for transformer-based long-term series forecasting,” *Artificial Intelligence Review*, vol. 58, no. 3, 2025, doi: 10.1007/s10462-024-11044-2.
[15] J. M. Oliveira and P. Ramos, “Evaluating the effectiveness of time series transformers for demand forecasting in retail,” *Mathematics*, vol. 12, no. 17, 2024, doi: 10.3390/math12172728.
[16] G. H. H. Nayak *et al.*, “Transformer-based deep learning architecture for time series forecasting,” *Software Impacts*, vol. 22, Nov. 2024, doi: 10.1016/j.simpa.2024.100716.
[17] E. R. H. P. Isaac and A. Sharma, “Adaptive thresholding heuristic for KPI anomaly detection,” *16th International Conference on COMmunication Systems and NETworkS*, pp. 737–741, 2024, doi: 10.1109/COMSNETS59351.2024.10427016.
[18] A. Samitas, E. Kampouris, and D. Kenourgios, “Machine learning as an early warning system to predict financial crisis,” *International Review of Financial Analysis*, vol. 71, 2020, doi: 10.1016/j.irfa.2020.101507.
[19] S. Puli, N. Thota, and A. C. V. Subrahmanyam, “Assessing machine learning techniques for predicting banking crises in India,” *Journal of Risk and Financial Management*, vol. 17, no. 4, 2024, doi: 10.3390/jrfm17040141.
[20] A. Eskandarany, “Adoption of artificial intelligence and machine learning in banking systems: a qualitative survey of board of directors,” *Frontiers in Artificial Intelligence*, vol. 7, 2024, doi: 10.3389/frai.2024.1440051.
[21] V. L. Heß and B. Damásio, “Machine learning in banking risk management: mapping a decade of evolution,” *International Journal of Information Management Data Insights*, vol. 5, no. 1, 2025, doi: 10.1016/j.jjimei.2025.100324.
[22] A. B.- García, A. Conde, U. Mori, and J. A. Lozano, “A review on outlier/anomaly detection in time series data,” *ACM Computing Surveys*, vol. 54, no. 3, 2022, doi: 10.1145/3444690.

[23] S. Crépey, N. Lehdili, N. Madhar, and M. Thomas, "Anomaly detection in financial time series by principal component analysis and neural networks," *Algorithms*, vol. 15, no. 10, 2022, doi: 10.3390/a15100385.

[24] S. O. Pinto and V. A. Sobreiro, "Literature review: anomaly detection approaches on digital business financial systems," *Digital Business*, vol. 2, no. 2, 2022, doi: 10.1016/j.digbus.2022.100038.

[25] B. Dumitrescu, A. Baltoiu, and S. Budulan, "Anomaly detection in graphs of bank transactions for anti money laundering applications," *IEEE Access*, vol. 10, pp. 47699–47714, 2022, doi: 10.1109/ACCESS.2022.3170467.

[26] P. Vanini, S. Rossi, E. Zvizdic, and T. Domenig, "Online payment fraud: from anomaly detection to risk management," *Financial Innovation*, vol. 9, no. 1, 2023, doi: 10.1186/s40854-023-00470-w.

[27] A. Song, E. Seo, and H. Kim, "Anomaly VAE-transformer: a deep learning approach for anomaly detection in decentralized finance," *IEEE Access*, vol. 11, pp. 98115–98131, 2023, doi: 10.1109/ACCESS.2023.3313448.

[28] P. Jathade, P. Besekar, and M. Sadu, "Currency conversion with real time API integration," *International Journal of Advanced Research in Science, Communication and Technology*, vol. 4, no. 1, pp. 410–416, 2021, doi: 10.48175/568.

[29] F. R. Alharbi and D. Csala, "A seasonal autoregressive integrated moving average with exogenous factors (SARIMAX) forecasting model-based time series approach," *Inventions*, vol. 7, no. 4, 2022, doi: 10.3390/inventions7040094.

[30] T. Xie and J. Ding, "Forecasting with multiple seasonality," *2020 IEEE International Conference on Big Data (Big Data)*, pp. 240–245, 2020, doi: 10.1109/BigData50022.2020.9378072.

[31] O. Triebe, H. Hewamalage, P. Pilyugina, N. Laptev, C. Bergmeir, and R. Rajagopal, "NeuralProphet: explainable forecasting at scale," *arXiv:2111.15397*, 2021.

BIOGRAPHIES OF AUTHORS

Yara Harb received her B.S. degree in Computer Science from the Lebanese American University (LAU), Beirut, Lebanon, in 2023, where she is currently pursuing her M.S. degree in Computer Science. Her research interests include data analytics, time-series forecasting, and financial technology applications. She has experience in Python, SQL, and Power BI for developing data-driven solutions and analytical dashboards. She completed an internship as a Data Analyst at Al Yafi IP Group, where she focuses on transforming complex datasets into actionable business insights. She was awarded first place in the best student project award at the 7th ICMMAAC-24 Conference in 2024. She can be contacted at email: yara.harb01@lau.edu.

Wissam Baaklini received his bachelor's degree in computer science with distinction from the Lebanese American University, where he is currently pursuing his master's degree. Dedicated to the fields of systems engineering, automation, and intelligent infrastructure, he has professional experience as a Core Systems and Infrastructure Engineer at BankMed, managing enterprise Linux environments, CI/CD pipelines, and Oracle middleware platforms. He previously served as a DevOps Engineer at Murex and as a main instructor for the Algorithms and Data Structures Lab at LAU. He holds certifications from AWS, NVIDIA, and MICRO1 and remains actively engaged in academic research, teaching, and the advancement of secure, scalable computing systems. He can be contacted at email: wissam.baaklini@lau.edu.

Dr. Nadine Abbas is an Assistant Professor of Computer Science at the Lebanese American University (LAU). She earned a BE degree in Computer and Communication Engineering from Notre Dame University (NDU) in 2008. She received her M.E. and Ph.D. degrees in Electrical and Computer Engineering from the American University of Beirut (AUB) in 2011 and 2017, respectively. Her research focuses on mobile edge computing and wireless communications, developing innovative solutions for next-generation heterogeneous networks, including 5G/6G mobile systems, the internet of things, and UAV-aided communications. She can be contacted at email: nadine.abbas@lau.edu.lb.

Prof. Seifedine Kadry received the bachelor's degree from Lebanese University, Beirut, Lebanon, in 1999, the dual M.S. degree from Reims University, Reims, France, and EPFL, Lausanne, Switzerland, in 2002, the Ph.D. degree from Blaise Pascal University, Aubière, France, in 2007, and the H.D.R. degree from Rouen University, Rouen, France, in 2017. At present his research interests include data science, education using technology, system prognostics, stochastic systems, and applied mathematics. He is currently the ABET program evaluator for computing, and the ABET program evaluator for engineering tech. He is also a fellow of IET, IETE, and IACSIT. He is also a distinguish speaker of the IEEE Computer Society. He can be contacted at email: skadry@gmail.com.