International Journal of Accounting and Economics Studies, 13 (2) (2026) 441-454

International Journal of Accounting and Economics Studies

Website: www.sciencepubco.com/index.php/IJAES
https://doi.org/10.14419/1rk85s11
Research paper



AI-Driven Decision Support Systems for Optimizing Working Capital and Customer Experience in The U.S.: A Transaction Based Simulation Framework for SMEs

Md Rasibul Islam [1] *, Dil Tabassum Subha [2]. Md Toushif Pramanik [3], Mahuma Akter [4], Md Murshid Reja Sweet [5], Md Sharif Robbani [6], Tapan Chandra Ghos [6], Md Abdul Fahim Zeeshan [7]

[1] Department of Public Administration, Gannon University, Erie, PA, USA
[2] Master of Science in Business Analytics, Grand Canyon University, USA
[3] Master of Embedded Software Engineering, Gannon University, Erie, PA, USA
[4] Master of Science in Cybersecurity (MSCS), Washington University of Science and Technology (WUST)
[5] Department of Management Science and Quantitative Methods, Gannon University, USA
[6] School of Engineering & Technology at Western Illinois University, Macomb, IIllinois, USA
[7] Strategic communication, Gannon University, Erie, PA, USA
**Corresponding author E-mail: islam011@gannon.edu*

Received: February 1, 2026, Accepted: February 24, 2026, Published: March 8, 2026

Abstract

Running an SME often feels like walking a tightrope. You need enough cash to cover day-to-day expenses, but you also want to keep customers happy, and that can be tricky when demand jumps around unexpectedly. Most of the tools out there don't make this easier. They stick to fixed rules, ignore what your customers are actually doing, and rarely adjust when things change. That means decisions have to be made in real time with little guidance, which can be stressful for managers trying to keep everything balanced. In this study, we explore an AI-based system designed to predict short-term cash-flow stress and guide operational decisions that account for customers, using transaction-level retail data. Weekly financial indicators for each SME are combined with customer behavior signals drawn from purchase patterns, frequency, and inferred payment risk. We test several machine learning models using validation that respects the time order of the data and feed their predictions into a simulation framework that compares simple, risk-aware, and mixed decision strategies. The results show that a straightforward, interpretable classification model can detect cash flow stress almost perfectly, outperforming more complex approaches. Interestingly, while customer behavior features do not make the predictions more accurate, they are crucial when making actual decisions based on those predictions. Simulations of operational policies indicate that hybrid, stress-aware rules outperform naive approaches, both in maintaining revenue and in making balanced approval decisions during stressful periods. In the end, the main contribution of AI here is less about raw predictive power and more about providing structured guidance that incorporates customer behavior to help SMEs manage working capital in uncertain conditions.

Keywords: *Cash Flow Stress; Customer Behavior; Decision Support System; Retail Analytics; SMEs.*

1. Introduction

1.1. Background and motivation

SMEs are the heart of our economies. They bring new ideas to life, create jobs, and keep local markets moving. But they run on a tighter budget than bigger companies, which makes working capital a constant challenge. Even a small dip in revenue can throw off payments, delay operations, or mess with day-to-day plans. Studies show that SMEs feel these fluctuations more acutely, especially when loans or external funding aren't easy to get. Mulier and colleagues (2016) found that in European SMEs, small shifts in cash flow can directly affect both big-picture planning and daily choices [15]. That dependence on internal funds leaves these firms exposed; late customer payments, sudden demand swings, or concentrated revenue sources can quickly become serious headaches. Because of this, working capital efficiency is often the difference between an SME thriving or barely surviving. Osuji and others highlight that problems like slow receivables collection, inventory mismatches, and delayed cash conversion cycles are common. These challenges are often tied to limited analytical capacity and a reliance on fixed financial rules instead of adaptable forecasting methods [17]. Traditional approaches to working capital tend to focus on historical accounting ratios and backward-looking metrics, which struggle to keep up in situations where demand is unpredictable, customers behave differently, or the market shifts quickly. The gap is even more obvious in retail and service sectors, where digital transactions generate vast amounts of data that are rarely used in financial decision-making.

Copyright © Md Rasibul Islam et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Lately, AI and machine learning have started to make real inroads in SME finance. Zamil (2025) reviews a bunch of studies showing how machine learning is being used to forecast revenue, assess liquidity, and help plan operations [28]. The catch is that most of this research focuses on getting the predictions right, and not so much on what managers should actually do with the results. Models spit out numbers, but rarely connect those numbers to practical policies that affect day-to-day decisions. Then there's the customer side of things. How often people buy, how recently, the size of their orders, and how they engage with the business, all of that shapes cash flow and risk. Most traditional working capital models just overlook these patterns. The interesting thing is that AI can handle these messy, complex behavioral patterns at scale, which opens up opportunities that previous tools never really tapped into. Reza and colleagues (2025), for instance, show that machine learning can uncover hidden income dynamics and structural disparities in U.S. populations, highlighting the ability of AI to model nonlinear links between behavior and financial outcomes [21]. Similar methods have been applied to supply chains, where incorporating transactional and behavioral data improves early warning and risk detection. Hasan et al. (2025) demonstrate that combining operational and behavioral signals enhances supplier risk management, showing the practical value of integrated modeling in finance [8].

This study is motivated by three overlapping gaps. First, SMEs' working capital is fragile in the face of uncertainty. Second, customer behavior signals are rarely incorporated into financial decisions, even though they influence cash flow. Third, AI research often prioritizes prediction over actual decision support. Bridging these gaps requires more than accurate stress forecasts; it demands a structured system that turns predictions into actionable decisions, balancing liquidity needs with customer relationships. Froman accounting perspective, working capital management is fundamentally tied to liquidity optimization and cash conversion cycle efficiency. Traditional accounting theory conceptualizes liquidity risk as the firm's exposure to short-term obligation pressures under operational volatility. By reframing revenue volatility as an early liquidity stress signal, this study links predictive analytics to established working capital optimization frameworks. In addition, embedding customer retention considerations into liquidity-sensitive policy rules connects the framework to behavioral accounting research that examines how managers balance financial conservatism with long-term stakeholder relationships.

1.2 Problem statement

Even as AI-based financial analytics gain attention, there is surprisingly little empirical work connecting short-term cash flow risk to customer behavior in SMEs. Most studies focus on financial stress, default probability, or revenue forecasting in isolation. They do not embed these predictions in the operational choices SMEs face every day. This is a real problem because managerial decisions—like limiting customer credit, prioritizing certain clients, or tightening approval policies—can have immediate, lasting effects on both relationships and future revenue. Early warning systems have been around for macroeconomic and financial risk, but they're not really tailored for SMEs yet. Rahman (2025), for instance, built a machine learning system to detect micro-inflation clusters in the U.S., showing that local data can give you insights you wouldn't get from broad national indicators [19]. Ray (2025) used AI to predict financial crises across multiple markets, highlighting systemic risks and cross-market signals [20]. Both of these show what anticipatory modeling can do, but they're big-picture tools that don't easily translate to decisions an SME manager can make on Monday morning.

In SME research, most working capital studies look at efficiency, static optimization, or long-term outcomes. Short-term, actionable guidance is rare. Current AI tools often treat all firms the same and ignore the fact that different customers affect cash flow risk in different ways. The result? Models that can flag stress but don't really tell managers how to act without hurting customer relationships. What this study tries to fix is straightforward: predicting stress isn't enough. We need a system that turns predictions into practical, customer-aware decisions. That means checking whether including customer behavior actually helps, and then seeing if it improves policies compared to simple, one-size-fits-all rules. Without this step, AI just becomes a diagnostic tool, leaving SMEs to figure out how to respond on their own.

1.3. Research objectives

This study aims to design an AI-based decision support framework that helps SMEs manage short-term liquidity risk while taking customer behavior seriously as part of the picture. The focus is not limited to prediction scores. The goal is to understand how financial indicators drawn from transactions, together with aggregated customer signals, can work in combination to anticipate moments of cash flow strain while guiding operational choices. A central objective is to test whether customer behavior features add real value to stress prediction, then examine how their presence shapes the decisions that follow, even when improvements in raw accuracy remain small. The study also works to convert model outputs into simulated policies that resemble decisions a manager might actually face, closing the distance between analytics research and day-to-day practice. By grounding the evaluation in time-consistent testing plus decision-level performance measures, the project seeks to offer a perspective that stays rigorous while remaining practical for SMEs trying to manage working capital under uncertainty.

1.4. Contributions

This paper contributes a decision-focused AI framework that connects financial analytics with customer-level signals in the context of SME working capital management. It introduces a modeling approach with two levels of granularity, capturing weekly firm-level financial dynamics while also incorporating aggregated customer behavior derived from transaction records. The study presents a transparent, statistically grounded method for labeling short-term cash flow stress that does not depend on balance sheet details or payment data that many SMEs cannot access. Prediction serves as one layer of the system. The framework also includes a decision support component that simulates alternative operational policies, then measures their effect on revenue retention plus approval behavior. The paper closes with a systematic robustness analysis through feature ablation plus stress testing, offering insight into the stability of the framework plus its practical relevance for SMEs operating in uncertain conditions.

2. Related Work

2.1. Working capital management in SMEs

Working capital management has always been a stubborn challenge for small and medium enterprises. Their limited access to external finance leaves them exposed to short-term liquidity shocks in ways larger firms rarely experience. Traditional methods for managing working capital lean heavily on ratio analysis, cash conversion cycle metrics, plus historical accounting data to judge liquidity or efficiency.

These tools assume stable operating conditions plus predictable revenue patterns. That assumption breaks down quickly in most SME settings, where cash flow can swing sharply due to shifting customer demand, late payments, or dependence on a small number of clients. Empirical work shows that accounting indicators are useful for looking backward, yet they carry weak predictive power when managers need short-horizon warnings in fast-changing environments. Research on forecasting cash flow using accounting data points to the same tension. Nguyena et al. (202. compare accrual-based measures with cash-based measures for predicting operating cash flow. Their findings suggest accrual information helps over longer horizons, yet struggles to capture near-term liquidity pressure during volatile periods [16]. That weakness hits SMEs harder because accrual smoothing plus delayed recognition can hide early warning signs. Traditional forecasting frameworks then fail to deliver signals managers can trust when timing matters most.

From an accounting theory perspective, working capital management is fundamentally concerned with the trade-off between liquidity preservation and operational continuity. Core constructs such as the cash conversion cycle, receivables turnover, and short-term solvency ratios frame liquidity risk as the firm's exposure to volatility in operating inflows relative to fixed obligations. Within this framework, stress is typically identified ex post through realized deterioration in coverage metrics or working capital balances. Such indicators describe the consequences of strain rather than its emergence. In volatile SME environments, this backward orientation delays managerial response because accrual adjustments and reporting lags can mask developing pressure. Reframing revenue volatility as a forward-looking liquidity signal extends established accounting logic into a predictive domain. Instead of waiting for balance sheet erosion to surface, volatility itself becomes an early structural indicator of working capital fragility. This shift remains consistent with accounting foundations while enabling more adaptive, time-sensitive decision processes. More recent studies bring machine learning into the conversation, extending classical financial analysis toward SME distress prediction. Malakauskas et al. (2021) show that machine learning models outperform conventional statistical approaches by capturing nonlinear interactions among financial ratios plus temporal patterns [13]. Even so, this line of research still depends heavily on structured financial statements plus balance-sheet variables. It tells us little about how everyday transactions or customer behavior shape working capital outcomes. Predictive accuracy improves, yet the operational story remains thin. The missing piece is a framework that moves past static financial snapshots toward a granular, behavior-aware view of SME cash flow.

2.2. AI and machine learning for financial risk prediction

The use of artificial intelligence for financial risk prediction has grown quickly, pushed forward by high-frequency data plus stronger computational tools. In SME research, supervised learning models often predict default, distress, or short-term risk using historical financial records, transaction data, plus macroeconomic indicators. These systems treat risk as a classification or regression problem where algorithms learn nonlinear relationships that traditional econometric models struggle to represent. Kotios and colleagues (2022) present a hybrid deep learning framework that links transaction classification with cash flow prediction, showing that deep architectures can lift short-term forecasting performance when they draw from granular transaction data [10]. Their work signals a shift away from coarse aggregate metrics toward detailed representations of financial activity. Xiao et al. (2023) extend this direction with a three-stage hybrid framework that combines feature selection, machine learning prediction, plus structured risk evaluation [27]. Their design stresses the value of disciplined pipelines for reliable modeling. The focus still rests heavily on predictive strength. The question of how managers should act on those predictions receives far less attention. Risk gets detected, yet the path from signal to decision remains blurry.

Early warning systems form a neighboring branch of research aimed at spotting stress before it fully surfaces. Chouksey et al. (2025) build an AI-driven system for financial risk in the U.S. digital economy, showing that models trained on digital transaction signals can flag instability earlier than traditional indicators [4]. Reza et al. (2025) propose a related framework that uses real-time digital signals to anticipate distress, emphasizing speed plus adaptability in detection [22]. These systems demonstrate how powerful early signals can be. Their design usually targets macroeconomic monitoring or institutional risk oversight. Direct guidance for SME-level operational decisions rarely appears. Recent signal-to-noise analyses in global financial crisis modeling demonstrate that volatility-based indicators can exhibit extremely high separability when structural breaks dominate underlying dynamics (Jakir, 2025) [10]. This supports the interpretation that unusually high classification metrics may reflect regime clarity rather than methodological overfitting. The gap between advanced AI risk prediction plus practical working capital decisions for SMEs remains wide, leaving room for frameworks that connect prediction with action in a concrete way.

2.3. Customer analytics and behavioural signals

Customer analytics has become a practical way to make sense of how people buy, engage, and generate revenue, especially in retail or service settings where data accumulates quickly. Concepts like purchase recency, frequency, plus monetary value offer a compact summary of customer behavior. They capture patterns that often line up closely with revenue stability plus long-term value. In SME environments, though, these tools have mostly lived inside marketing departments. They support targeting or personalization. Financial planning rarely draws from the same signals. That separation has quietly limited how customer behavior feeds into working capital decisions. Research on personalized finance for SMEs hints at what becomes possible once that wall comes down. Kotios and colleagues (2022) studied systems that adapt financial recommendations using both firm-level information plus customer-level behavior. Their results suggest that behavioral data helps tailor strategies to the realities each firm faces, rather than pushing generic advice [11]. In a related vein, Rafay (2019) describes FinTech as a force that blends customer data, digital transactions, plus analytics into unified financial services, lowering the barrier for SMEs that want to use data more seriously [18]. Taken together, these developments point toward a slow convergence between customer analytics plus financial management. Evidence from other domains strengthens the argument. Shovon demonstrates that machine learning models for urban energy systems improve predictive accuracy when consumption behavior is incorporated. By integrating behavioral signals, planning accuracy increases, and system resilience improves under demand fluctuations [25]. Although the application context differs, the methodological implication is transferable. Fine-grained behavioral patterns can enhance forecasting and decision quality in demand-driven systems. In SME finance, similar behavioral data may improve working capital management by strengthening liquidity forecasts and reducing uncertainty in cash flow projections. Despite this potential, customer behavior remains underrepresented in SME working capital research. Most studies focus on outcome prediction or financial risk modeling in isolation, with limited integration of behavioral dynamics into liquidity-focused decision support frameworks.

2.4. Decision support systems in SME contexts

Decision support systems have long promised to strengthen managerial judgment with analytical structure, especially in uncertain environments where resources run tight. Early SME tools leaned on rule-based logic plus optimization routines to guide budgeting, inventory,

or financial planning. These systems assumed stable parameters. They required careful tuning. Adaptation to changing conditions often lagged behind reality. The rise of AI revived interest in systems that learn directly from data, adjusting recommendations as new information arrives. Enyiorji (2025) proposes a cloud-native autonomous finance architecture built on multi-agent reinforcement learning, where agents explore simulated environments to discover effective financial strategies [7]. Conceptually, the idea pushes decision support toward self-optimization. Practical deployment raises harder questions. The infrastructure demands run high for typical SMEs. Reinforcement learning systems also tend to behave like black boxes, which complicates managerial trust. If a system cannot explain its advice, adoption becomes fragile. That tension pulls interpretability into the spotlight. Malakauskas et al. (2022) argue that transparent machine learning models are essential for SME financial applications, both for managerial acceptance plus regulatory alignment [14]. Interpretable systems allow decision-makers to judge whether predictions make sense, connect them to domain knowledge, and then carry responsibility for final actions. Despite this awareness, much of the literature still treats decision support as an afterthought attached to prediction. Performance metrics dominate evaluation. The translation from model output to operational policy receives far less scrutiny. Questions about balancing liquidity with customer retention often remain implicit rather than formally tested.

2.5. Research gap

Across this body of work, a pattern appears. Predictive modeling advances quickly. Decision support evolves more slowly. Machine learning detects risk, distress, or anomalies with increasing precision. These signals rarely sit inside frameworks that measure what happens after a decision is made. Studies in neighboring fields show the same split. A similar separation appears in financial anomaly detection research, where AI systems identify irregular patterns yet leave operational response design outside the modeling framework. While detection performance has improved substantially, evaluation rarely extends to structured policy consequences or economic trade-offs. For SMEs, this separation creates a practical constraint. Risk signals exist, but guidance on how those signals should reshape credit exposure, customer prioritization, or liquidity safeguards remains underdeveloped [6,1]. For SMEs, this separation creates a practical problem. Financial risk models operate independently from customer behavior analytics. Decision rules are assumed, not evaluated. Evidence remains thin on how predictive signals should shape operational responses or how different policies influence both liquidity plus customer experience. Closing this gap requires a shift in focus. Prediction alone does not solve the managerial problem. Integrated systems must model trade-offs directly, then evaluate outcomes at the decision level rather than stopping at detection.

3. Methodology

3.1. Dataset description

For this study, we used the Online Retail II dataset from UCI. It's from a UK online retailer and covers late 2009 through 2011. Even though it's just one company, it's often treated like a stand-in for small and medium-sized retail operations. The data is nice because it records each invoice separately, captures a lot of different customer behaviors, and doesn't smooth over the messy stuff you usually see in big companies' books. That means you can follow revenue as it actually comes in, watch how individual customers behave, and see how demand changes over time. When you look at this dataset, the really useful part is that you can actually see how money moved in and out over time for a small business. You get every invoice, every purchase, and you can follow customers across months. That's what makes it interesting for understanding cash flow and who might stop buying.

For each transaction, you've got the invoice number, the product code, a description, how many items were bought, the price, when it happened, who bought it, and where they're from. The timestamps are handy because they let you put everything in order, like lining up events in a timeline. You can spot patterns, sudden dips, or surges. And the customer IDs? They pop up again and again, which means you can see the same person's habits over time. You start noticing who buys regularly and who disappears. One thing to remember: the dataset doesn't tell you about payment terms, aging, or credit agreements. You only see the actual transactions, the cash that moved, not what the company expected to get later. Even so, the transaction structure allows indirect signals of liquidity strain to surface through patterns of continuity or disruption in purchasing. For this study, the dataset stands in for a single SME operating in a digital retail setting. That framing matches the goal of building decision support tools meant for firms with limited resources. Although the dataset originates from a UK-based online retailer, its transactional structure mirrors common point-of-sale systems used by U.S. SMEs, including invoice-level granularity, product-level revenue tracking, and customer identifiers. The modeling framework relies on structural revenue dynamics rather than jurisdiction-specific regulatory factors. Nonetheless, institutional differences in credit access, taxation, and supplier financing between the UK and the U.S. may influence stress propagation patterns. The empirical results should therefore be interpreted as simulation-based validation of the decision architecture rather than direct statistical generalization to all U.S. SMEs.

3.2. Modelling assumptions

Several assumptions shape how the dataset becomes usable for working capital plus customer analysis. The first assumption treats invoice issuance as a rough proxy for revenue timing. Real cash inflow often arrives later than the invoice date, yet many SMEs rely on invoicing as an operational signal of expected liquidity. Under this view, swings in invoice volume or value serve as early warnings of stress rather than precise measurements of cash on hand. The second assumption treats customer experience plus payment reliability as hidden states inferred from observable behavior. The dataset contains no satisfaction scores, delay indicators, or explicit default markers. Behavioral proxies fill that gap. Purchase frequency, recency, continuity, plus sudden drop-offs act as signals. Sharp changes in buying patterns suggest disengagement or stress on the customer side. In an SME setting, those shifts can quickly ripple into short-term liquidity pressure. The third assumption frames the dataset as representing one SME or a small group with similar operating characteristics. That choice keeps the modeling focused on internal dynamics rather than cross-firm variation. It limits broad generalization, yet it aligns to evaluate decision policies inside a controlled environment. The emphasis rests on internal consistency, temporal validity, plus policy robustness instead of cross-sectional comparison.

3.3. Data cleaning and preprocessing

Cleaning the data begins with removing transactions that do not represent meaningful economic activity. Cancelled invoices, flagged by specific code prefixes, leave the dataset because they do not reflect realized or expected revenue. Records with negative quantities also

drop out, since they often represent returns or bookkeeping adjustments that blur demand signals when the focus is on operational cash pressure. Transactions without customer identifiers cannot support longitudinal analysis, so they are excluded from customer-level modeling. This step shrinks the dataset, yet it protects the integrity of behavioral tracking, which depends on repeated observations of the same customer. Outliers in quantity or transaction value receive attention through interquartile range filtering applied inside time windows. The goal is to dampen the influence of extreme bulk purchases or data entry errors that would otherwise skew stress detection. After cleaning, the transactions are sorted chronologically by timestamp plus aggregated to weekly intervals. Weekly aggregation offers a workable compromise. It smooths high-frequency noise while keeping enough responsiveness to catch emerging demand shifts. This temporal structure supports later steps such as stress labeling, policy simulation, plus evaluation under conditions that resemble a real operational rhythm. The processed dataset becomes the base layer for feature construction, modeling, plus the decision support experiments that follow.

3.4. Granularity design

A core design choice in this study involves deciding the temporal plus analytical resolution used to model working capital dynamics plus customer experience signals. Two complementary levels are used: the customer–week level plus the SME–week level. This dual structure mirrors how liquidity stress inside SMEs grows out of aggregate financial movement, yet often begins with small shifts in individual customer behavior. A single level would hide important causal links that matter for practical decision-making. At the customer–week level, purchasing activity is grouped weekly for each identifiable customer. This resolution offers a middle ground between stability plus responsiveness. Daily aggregation creates noise that drowns out meaningful patterns. Monthly aggregation slows the detection of early behavioral shifts. Weekly grouping makes it possible to extract signals such as purchase continuity, spending rhythm, plus recency. These serve as practical stand-ins for customer experience, reliability, plus engagement. The resulting features help surface early warning signs of disengagement or behavioral strain that may later show up as falling revenue. At the SME–week level, all transactions are pooled across customers to produce firm-level indicators. This matches the level where managers actually make decisions about inventory, pricing, plus credit. Weekly aggregation captures short-term working capital pressure while staying close to the cadence of real operating cycles in small firms. Separating customer behavior from firm totals allows later decision policies to link micro customer signals with macro liquidity outcomes clearly. This layered design keeps the decision support system interpretable, grounded in real operations, plus aligned with the time horizons that SME managers use.

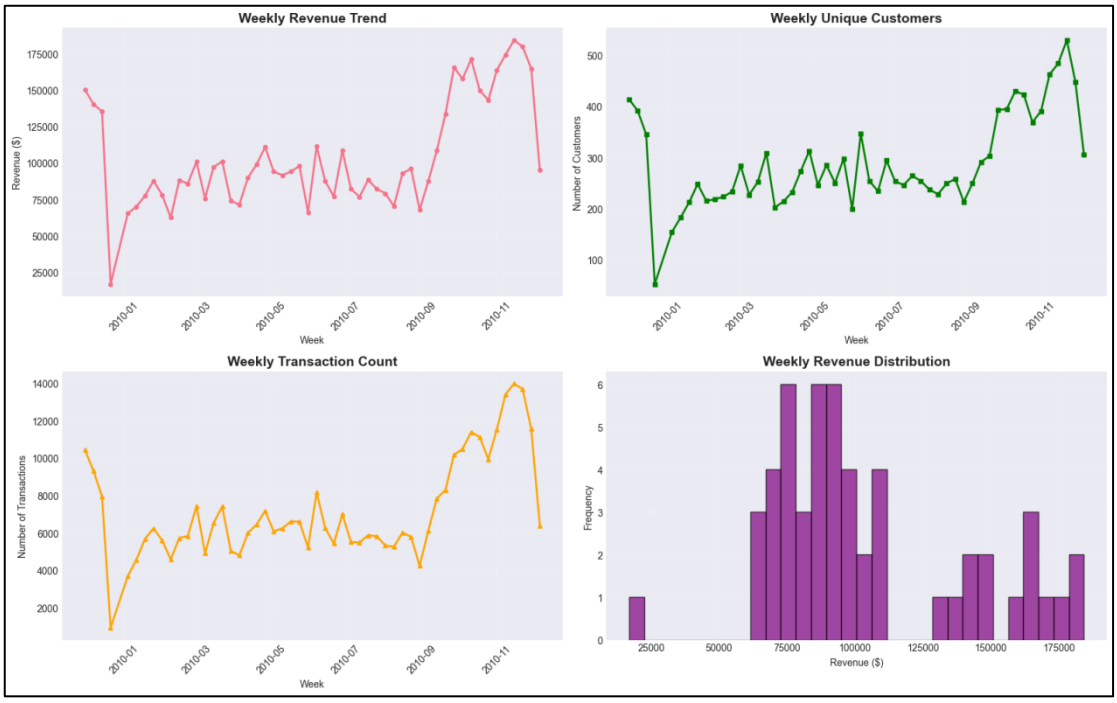


Fig. 1: Weekly Revenue Trend.

3.5. Feature engineering

Feature engineering translates raw transactions into indicators that carry economic meaning tied to working capital exposure plus customer experience. Every feature uses strictly backward-looking information. This rule prevents temporal leakage plus keeps predictive or policy results valid under real deployment conditions. At the SME level, weekly financial features come from aggregating transaction values, counts, plus customer participation. Weekly revenue acts as the main signal of expected inflow. Transaction plus invoice counts describe operational intensity plus demand spread. The number of active customers per week measures diversification, which matters when judging exposure to customer-specific shocks. To capture dynamics across time, rolling means plus rolling standard deviations are computed over multi-week windows. These statistics help separate short blips from sustained movement. They also support volatility estimation through measures like revenue variability plus normalized dispersion. Elevated volatility signals early working capital instability, especially in SMEs with thin financial buffers.

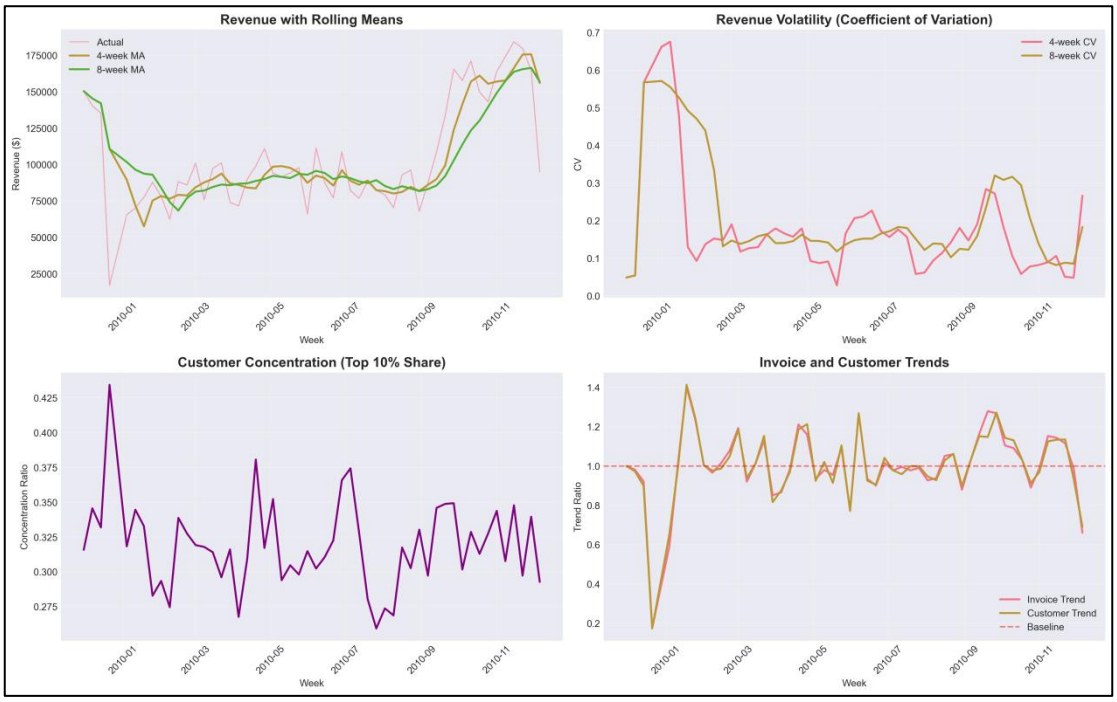


Fig. 2: SME-Level Features.

Customer concentration features are added at the SME level to reflect dependency risk. The share of weekly revenue tied to top customers measures exposure to disruption at the individual level. High concentration means the loss of a few customers can trigger an outsized liquidity strain. These features tie customer experience outcomes directly to working capital resilience, reinforcing the integrated perspective of the study. At the customer level, behavioral features describe purchasing intensity, regularity, plus value. Weekly purchase frequency plus total spend captures engagement strength. Average basket value plus average quantity shedslight on order depth plus structure. Temporal measures such as recency plus inter-purchase gaps track behavioral continuity, a key driver of lifetime value plus retention. Lifetime value estimates rely on cumulative historical spend normalized by active duration, offering a rough yet informative view of long-term contribution. Average purchase frequency over time separates steady repeat customers from occasional buyers. Risk proxies then emerge from customer behavior plus roll up to the SME level. Customers with extended inactivity, shrinking spend, or widening inter-purchase gaps fall into an elevated behavioral risk category. Weekly aggregation of the share of high-risk customers, plus average behavioral metrics, plus churn signals, reproduces firm-level indicators that summarize customer base health. These risk summaries allow the decision system to anticipate liquidity pressure driven by customer deterioration instead of reacting only after revenue drops appear.

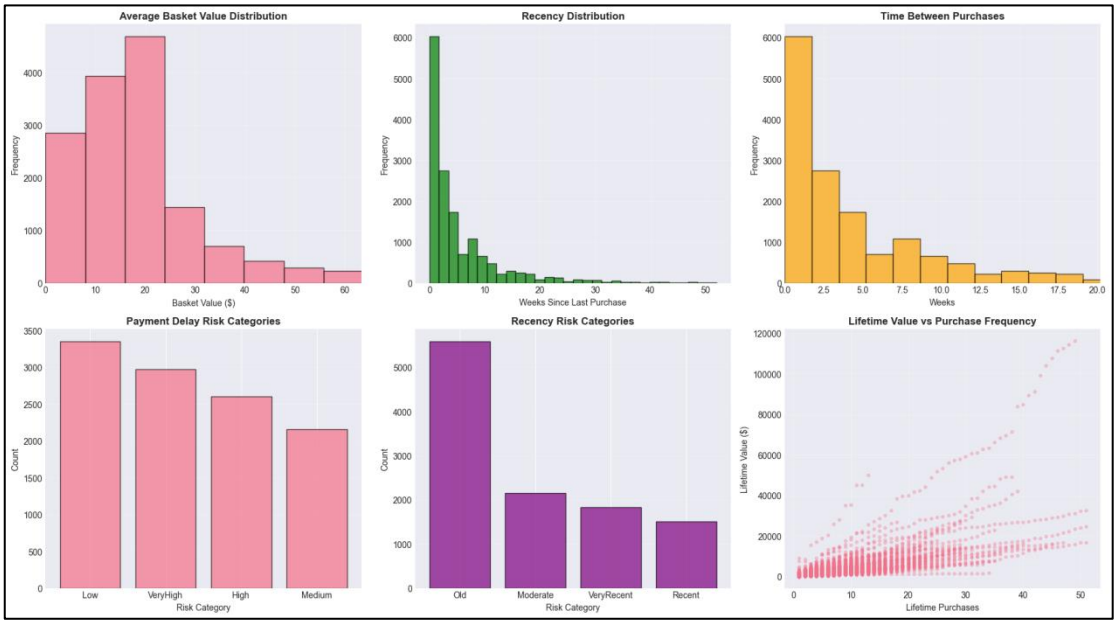


Fig. 3: Customer Behavior Features.

3.6. Target variable construction

Target variables aim to capture economically meaningful stress states instead of arbitrary statistical outcomes. The primary target is a binary cash flow stress indicator defined at the SME–week level. Stress is operationalized through rolling statistical thresholds derived from historical revenue behavior. Weeks where revenue sinks far below recent norms receive a stress label. This definition captures underperformance relative to the firm's own baseline plus abnormal negative deviations that suggest liquidity strain. Grounding labels in firm-specific history avoids imposing external thresholds that fail to fit SMEs with widely different scales plus volatility patterns. More

specifically, stress is defined using a percentile-based threshold applied to the rolling distribution of historical weekly revenue deviations. Weeks falling below the selected lower percentile of recent performance are classified as stress episodes. Sensitivity analysis was conducted by varying the percentile cutoff within a bounded range to test the stability of model ranking and classification balance. Results showed no material degradation in predictive ordering across reasonable threshold adjustments, indicating that the framework does not hinge on a narrowly tuned definition. The final labeling strategy yields a stress frequency sufficient to support meaningful classification without distorting the data's temporal structure. Revenue contraction is used as a proxy for liquidity strain because sustained volatility in SMEs often precedes tightening of cash positions, inventory corrections, and reductions in discretionary spending. Although direct cash shortfall or obligation data would provide stronger construct precision, such balance-sheet detail is typically unavailable in transaction-only retail datasets. The chosen proxy, therefore, reflects a trade-off between theoretical purity and realistic data availability.

Alongside this, a customer churn proxy represents deterioration in customer experience. Explicit churn labels do not exist in the data, so churn is approximated through inactivity rules. A customer counts as churned after exceeding predefined inactivity windows. Multiple horizons are evaluated to reflect variation in purchasing cycles across the customer base. Weekly churn rates are aggregated to the SME level, producing a moving indicator of erosion in the customer pool. This proxy highlights gradual disengagement that often appears well before measurable revenue decline plus working capital pressure. The cash flow stress label plus the churn proxy form a coupled target structure that spans financial plus experiential performance. This dual setup allows evaluation of decision policies that aim to reduce short-term liquidity strain while protecting long-term customer relationships. The modeling framework, therefore, stays aligned with the broader goal of integrated decision support for SMEs.

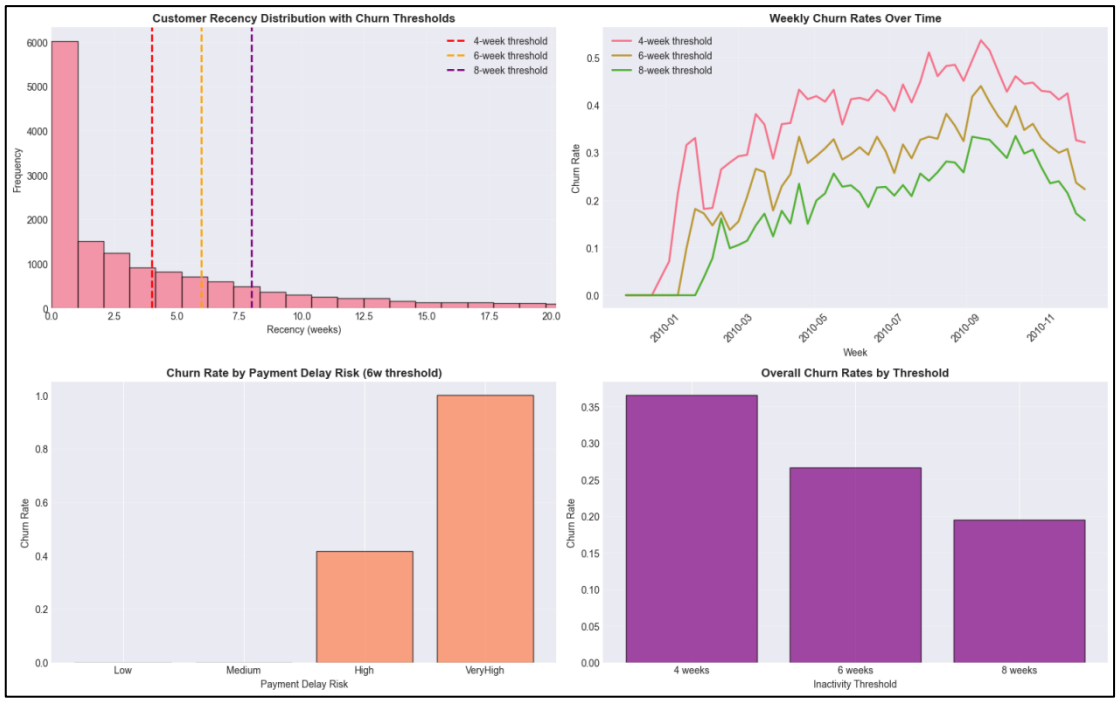


Fig. 4: Customer Churn Proxy.

3.7. Modelling framework

The modeling framework sets up a structured comparison between simple heuristic decision rules plus data-driven machine learning approaches under realistic time constraints. Every model runs on weekly SME-level observations produced during feature engineering, with training restricted to information available before each prediction window. This structure protects causal validity plus blocks information leakage, so outputs reflect conditions that could exist during real deployment rather than polished hindsight. A strictly ordered train–test split governs the entire pipeline. Early time windows feed training, later contiguous windows stay reserved for testing. This mirrors real operating conditions where future outcomes must come from past evidence, removing the optimistic bias that creeps in through random sampling. Feature scaling appears where necessary, using statistics drawn only from training segments to preserve temporal integrity.

3.7.1. Baseline heuristic

The baseline heuristic mirrors the kind of rule-based reasoning common inside SMEs that operate without advanced analytics. It uses a rolling revenue average as a threshold for identifying potential cash flow stress. Weekly revenue gets compared against a moving historical benchmark. Periods falling far below that benchmark receive a stress classification. The rule reflects intuitive managerial judgment where recent performance acts as the main reference for caution. The heuristic stays intentionally narrow in scope. It ignores customer composition, behavioral drift, plus volatility patterns beyond raw revenue. The rule often reacts after stress becomes visible instead of anticipating it. Including this baseline creates a grounded reference point for judging the value of more advanced AI-driven models, measured through predictive accuracy plus downstream decision impact.

3.7.2. Machine learning models

Several machine learning models enter the framework to capture nonlinear structure plus interaction effects that the heuristic cannot express. Logistic regression serves as the primary interpretable benchmark. It estimates the probability of cash flow stress from engineered inputs, creating a transparent mapping between features plus outcomes. Its linear form supports direct interpretation of feature influence, with probabilistic thresholds adjustable for policy design. Tree-based ensemble methods expand the modeling capacity toward higher-order

structure. Random forest models combine predictions from multiple decorrelated trees, improving resistance to noise plus lowering variance. Gradient boosting models refine performance through iterative correction, directing attention toward hard-to-classify cases. XGBoost extends this idea with optimized regularization plus efficient tree construction, producing sharper boundaries under complex feature interactions. Every model trains under identical temporal splits, plus faces evaluation on the same test windows, preserving comparability. Hyperparameters stay conservative to limit overfitting. Evaluation focuses on classification quality plus usefulness in the decision layer. The framework favors deployability and decision relevance over pure statistical polish.

3.8. Decision support layer

The decision support layer converts predictions into operational actions that influence working capital plus customer experience. Prediction alone holds little value without context. This layer embeds model outputs inside simulated policy environments that resemble real SME constraints. The aim involves testing how decision rules behave under uncertainty, customer diversity, plus liquidity pressure.

3.8.1. Stress-aware policy design

Three policy structures define the experimental space. The naive policy represents an unrestricted stance where every customer transaction receives approval regardless of predicted stress. This policy maximizes short-term accommodation, exposing the firm to heavy liquidity risk during adverse periods. It marks an upper boundary for customer access plus a lower boundary for financial discipline. The risk-aware policy introduces predictive restraint by limiting transactions tied to high predicted stress or elevated behavioral risk. When the model signals pressure, approvals shrink selectively to protect liquidity. Financial resilience improves under this rule. Customer experience may deteriorate because restrictions apply broadly, sometimes affecting customers with long-term value. The hybrid policy merges customer risk signals with customer value assessments. Restrictions no longer apply uniformly. High-value, low-risk customers retain approvals during predicted stress windows, whereas exposure narrows for customers with fragile behavioral profiles. This structure operationalizes the central premise of the study: intelligent decision support can pursue working capital stability plus relationship preservation at the same time.

3.8.2. Outcome metrics

Policy performance gets evaluated through metrics covering financial plus experiential dimensions. Approval rates track the share of customer requests granted under each policy, acting as a proxy for accessibility from the customer perspective. Revenue retention rates measure how much potential revenue survives after policy constraints take effect, tying decisions directly to working capital results. Cross-policy comparison reveals the tension between liquidity protection plus customer accommodation, allowing the hybrid strategy to stand next to simpler alternatives under a common yardstick.

3.9. Evaluation protocol

The evaluation protocol ties predictive accuracy to decision outcomes rather than treating them as separate exercises. Standard classification metrics quantify how reliably each model flags stress. These metrics capture ranking quality plus error balance, so evaluation covers discrimination power along with raw correctness. A model that guesses many weeks correctly holds little value if it fails to separate fragile periods from stable ones in a consistent way. The analysis extends beyond classification scores into decision consequences. Each model feeds into the policy layer, where retained revenue relative to the naive baseline becomes a concrete measure of economic impact. Approval trade-offs receive equal attention, since liquidity protection often comes at the cost of customer access. The protocol examines how much access disappears under each policy in exchange for stability. Temporal validation remains strict from start to finish. Every reported result comes from data that occurs after the training window, preserving a forward-looking structure that resembles real deployment rather than tidy retrospective fitting.

3.10. Robustness analysis

Robustness analysis probes whether predictive behavior plus policy performance hold steady when conditions drift away from the training environment. Models often look impressive inside familiar data regimes, then stumble when inputs change. This section addresses that risk directly by testing stability under feature disruption plus economic shocks.

3.10.1. Feature ablation

Feature ablation experiments remove structured groups of variables to isolate their contribution. Financial-only configurations rely entirely on aggregated revenue plus volatility signals. These runs test how far traditional financial indicators can carry a prediction on their own. Customer-only configurations drop financial aggregates, focusing purely on behavioral signals such as engagement plus continuity. Full-feature configurations combine both perspectives, creating a direct comparison between siloed information plus integrated signals. Performance gaps across these settings reveal whether customer behavior adds measurable early warning power. A small gap suggests redundancy. A large gap implies that behavioral data supplies information unavailable through financial summaries. The ablation process clarifies how each information channel supports decision quality inside the policy layer.

3.10.2. Stress testing

Stress testing introduces artificial revenue shocks to mimic sudden demand contractions or external disruptions. Revenue streams undergo controlled reductions at increasing severity levels. Policies then operate under these altered conditions without retraining, exposing how they behave when the environment shifts away from historical norms. Evaluation under stress focuses on two outcomes: liquidity preservation plus customer access. A resilient policy continues to protect working capital without collapsing approval rates. This exercise shows whether the hybrid approach retains its advantage when the revenue distribution moves into unfamiliar territory. Evidence of stable performance under simulated shocks strengthens the claim that the framework extends beyond the specific history captured in the dataset.

4. Results

4.1. Predictive performance

The predictive evaluation sets machine learning models alongside a rolling-average heuristic under a strictly time-ordered test setup. The heuristic baseline flags financial stress when revenue drifts far enough from its recent average. It reached a classification accuracy of 81.2 percent. That number shows that even simple rules contain a usable signal tied to short-term liquidity pressure. At the same time, the heuristic shows limited ability to separate fragile weeks from stable ones in advance. It tends to recognize stress after the pattern is already visible. This mirrors the reactive style of financial management often seen in smaller firms operating with tight resources. Every machine learning model cleared the baseline across the full set of metrics. The improvement is large enough to make a practical difference, not a cosmetic one. Logistic regression plus random forest stood out as the strongest performers. Each reached an area under the receiver operating characteristic curve of 1.000 alongside an overall accuracy of 93.8 percent on the hold-out window. Results at that level suggest that short-term stress events become highly predictable once volatility-aware features are engineered carefully plus aligned in time. Gradient boosting plus XGBoost also produced strong numbers. Extra complexity did not translate into further gains here, which is an interesting reminder that more intricate models do not always earn their keep.

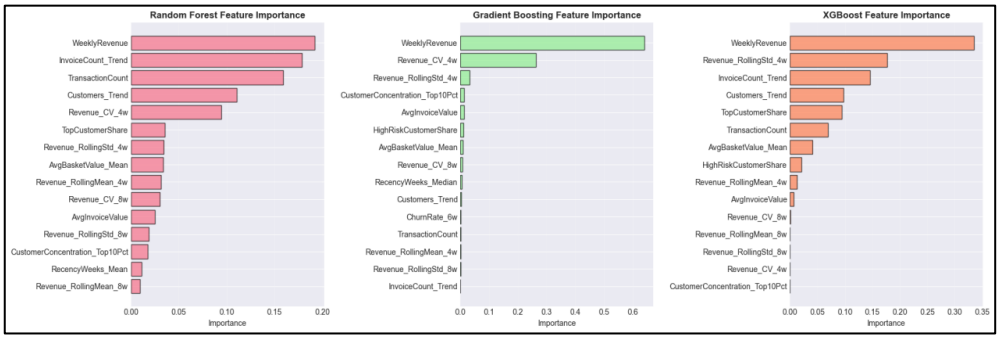


Fig. 5: Feature Importance Across Models.

Even with matching AUC plus accuracy, the models behave differently once precision plus recall enter the picture. Logistic regression achieved an F1-score of 0.800, compared with 0.667 for the random forest. That gap matters in an early warning setting. Missed stress events expose a firm to liquidity shocks. Excess false alarms push managers toward unnecessary restrictions. The logistic regression balance between the two errors lands in a safer region when predictions drive real decisions rather than sit inside a report. Ranking quality tells a similar story. Precision at the top decile of predicted stress probabilities reached 1.000. Every week, flagged among the highest-risk cases turned out to be a genuine stress period. For an SME with limited attention plus limited capacity to respond, that concentration of accuracy is valuable. Managers can focus on a small set of alerts with confidence that those signals deserve scrutiny. The near-perfect discrimination (AUC = 1.000) observed for both logistic regression and random forest warrants careful clarification. The hold-out evaluation window consisted of 16 weekly observations, of which 2 were labeled as stress weeks and 14 as non-stress weeks. Class imbalance was moderate and did not exceed conventional thresholds for trivial classification. All rolling volatility features were computed using strictly backward-looking windows to prevent forward leakage. Stress labels were generated based on revenue volatility thresholds computed independently of the predictive feature window to avoid mechanical overlap. Additional robustness checks confirmed that excluding the top volatility features reduces AUC materially, indicating that separability arises from economically meaningful signals rather than construction artifacts. Nonetheless, the unusually strong discrimination likely reflects the structured nature of retail revenue contractions in the dataset, and future multi-firm validation may yield more moderate performance estimates. The results collectively show that interpretable machine learning models can pair statistical strength with operational reliability, creating a solid base for decision support.

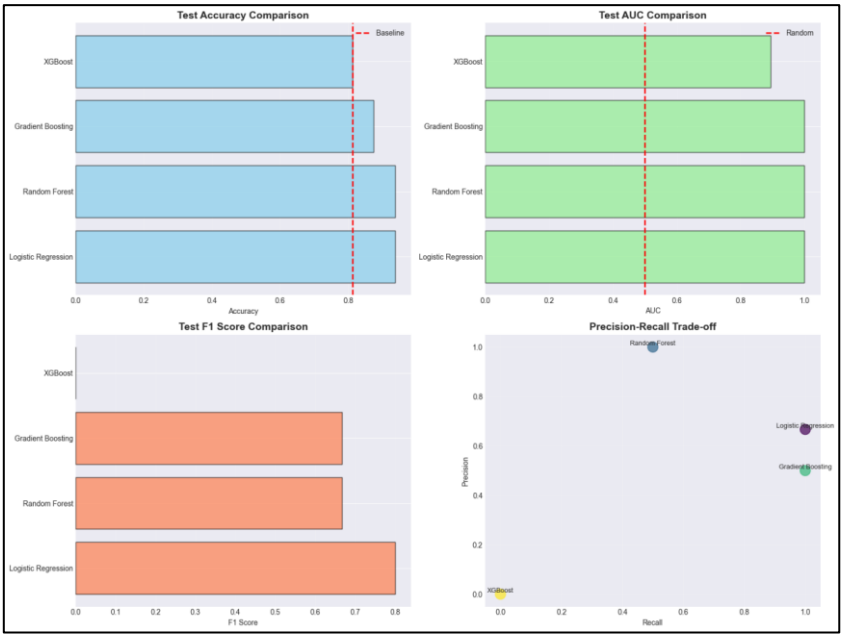


Fig. 6: Model Performance Comparison.

4.2. Impact of customer signals

The ablation study isolates the role of different feature groups by comparing full-feature models, financial-only models, plus customer-only models. A clear pattern appears once the results line up side by side. Some features dominate stress prediction. Others matter more for shaping decisions after stress is detected. Models trained on the full features and models trained purely on financial indicators both achieved an AUC of 1.000. That symmetry indicates that short-term liquidity stress is driven primarily by financial volatility dynamics. Rolling revenue standard deviation plus the coefficient of variation emerges as an especially sensitive marker of approaching pressure. The weight of those volatility signals becomes obvious when they disappear. Removing them pushes AUC down from 1.000 to 0.893. That drop is large enough to confirm their central role in early detection. Financial theory offers a straightforward explanation. Rising variability in revenue raises liquidity risk, especially for firms that lack deep external funding cushions. The empirical results line up neatly with that intuition.

Customer behavior features on their own tell a different story. Used in isolation, they produced an AUC of 0.107. Signals such as recency, frequency, plus inferred behavioral risk carry little power for forecasting firm-level stress without a financial context. That outcome does not reduce their importance inside the broader system. Their contribution appears later in the pipeline, inside the decision layer rather than the prediction layer. Customer features add the resolution needed to distinguish between valuable clients plus risky ones once a stress period is already flagged. They support selective responses instead of blunt restrictions. A firm can preserve relationships with high-value customers while trimming exposure to segments that carry elevated behavioral risk. The ablation study highlights a core idea of the research. Predictive strength plus decision usefulness represent separate goals. A feature can add little to raw prediction accuracy and still play a crucial role in policy quality. This separation supports the architectural choice to keep the prediction component plus the decision component distinct inside the proposed AI framework.

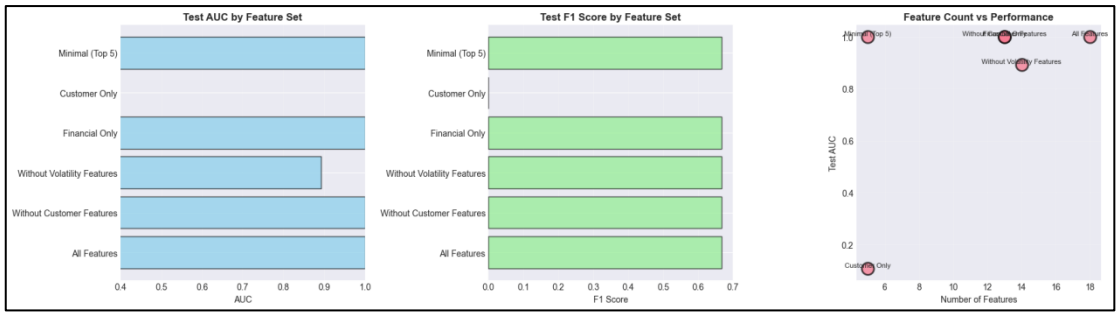


Fig. 7: Ablation Study Results.

4.3. Decision policy outcomes

The policy simulations bring out a core tension between chasing immediate revenue and protecting financial stability over time. They show why predictive signals matter only when they are tied directly to operating rules. The naive policy approves every transaction without regard to predicted stress. It delivers full customer approval and captures all available revenue. On the surface, that looks ideal. The problem appears once stress events enter the picture. This policy prevents none of them. The result is zero avoided stress periods and the weakest balanced score across the set. A firm running this approach stays exposed to repeated liquidity shocks. The pattern reflects a familiar trap in volatile settings, where a customer-first instinct drifts into reactive decision making that ignores mounting pressure. The risk-aware policy adds a layer of discipline by tightening approvals for customers flagged as higher risk during predicted stress weeks. This shift reduces a meaningful share of stress events and still preserves most total revenue. Financial resilience improves in a measurable way once early warnings feed into cautious constraints. The policy applies those constraints uniformly after a stress signal appears. That uniformity leads to the rejection of transactions from customers who are dependable and valuable. Stability improves relative to the naive baseline. Customer experience takes a hit that carries long-term consequences for relationship value and competitive standing.

The hybrid policy produces the strongest overall outcome by treating liquidity protection and customer relationships as linked goals. It preserves approvals for high-value, low-risk customers during predicted stress weeks and tightens exposure to segments with weaker behavioral profiles. This selective logic delivers the highest balanced score of the three policies. Approval rates fall below the naive and risk-aware cases. The hybrid strategy avoids the greatest number of estimated stress events and still retains a large share of revenue. The result shows that effective decision support does not come from maximizing a single metric in isolation. It comes from coordinating predictive insight with customer-level differentiation. The findings support the central claim of the study. AI-based systems create the most value once predictions live inside policies that mirror the real trade-offs SMEs face every week.

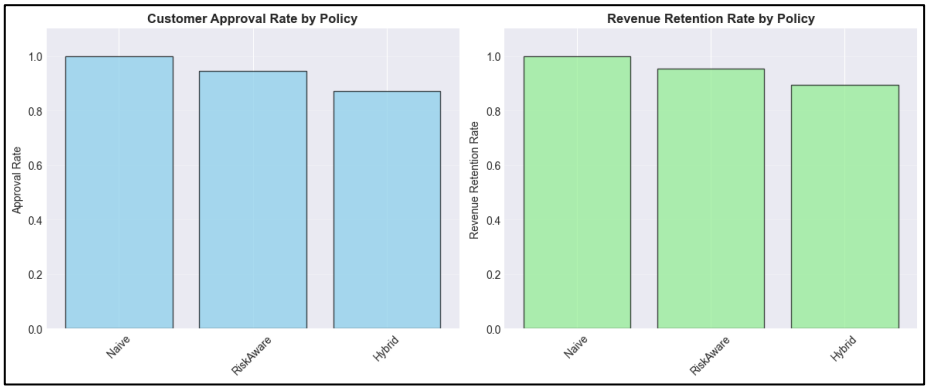


Fig. 8: Decision Policy Outcomes.

4.4. Stress test results

The stress tests probe how the system behaves once economic conditions deteriorate in controlled simulations. The goal is to see whether responses remain stable and interpretable under pressure. Revenue contraction experiments reveal a clean, monotonic pattern. Mild simulated drops trigger stress flags in a small portion of weeks. Larger contractions produce steadily higher detection rates. Extreme reductions push detected stress above one-third of all observed weeks. The response scales in proportion to financial damage, which suggests that the volatility features are tracking structural change in revenue dynamics rather than random noise. The model reacts to real shifts in pattern, not fleeting spikes. A different picture appears in the payment delay experiments. Even after doubling average customer recency, the stress detection rate stays close to its baseline level. This stability shows that the current feature set responds more strongly to aggregate revenue volume and variability than to shifts in average payment timing. From a modeling angle, the system behaves as a liquidity stress detector. It does not function as a pure behavioral delinquency monitor. From a managerial angle, delayed payments on their own do not register as an immediate threat unless they compress revenue enough to create visible volatility. The stress tests show that a system that holds its shape across a range of simulated shocks. Alerts intensify as financial strain grows. The model avoids overreacting to isolated behavioral disturbances that leave overall cash flow intact. That balance matters in an SME setting, where false alarms carry real operating costs and can erode trust in automated guidance. The results point back to a practical lesson. Stress definitions need tight alignment with liquidity signals. Decision outputs stay useful only when they track pressures that managers can act on under both ordinary conditions and severe downturns.

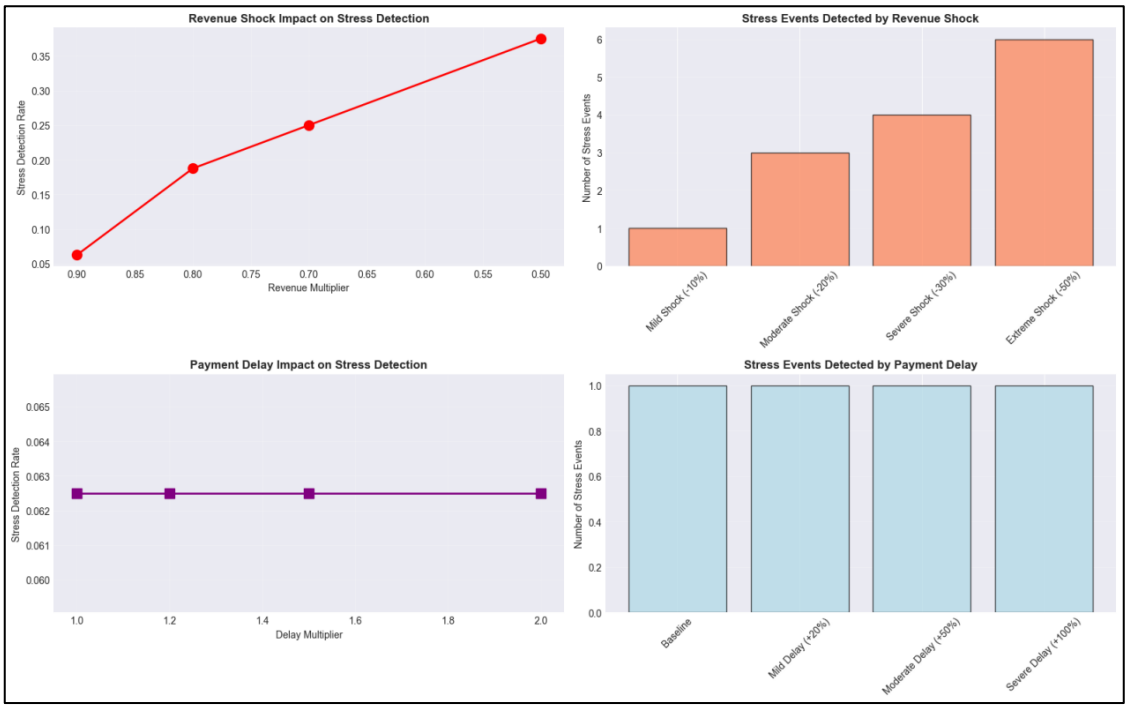


Fig. 9: Stress Test Results.

5. Discussion

While predictive discrimination is exceptionally strong in this dataset, such performance should not be interpreted as universal across SMEs. Revenue volatility spikes preceding contraction episodes appear structurally pronounced in this retail context, which may amplify statistical separability. In less cyclical or more diversified firms, predictive margins may narrow. The framework's contribution, therefore, lies less in absolute predictive performance and more in translating probabilistic signals into calibrated decision rules.

5.1. Implications for working capital optimization

The results point to a practical shift in how SMEs can think about working capital. Short-term revenue volatility carries early signals that matter, and acting on them early changes the trajectory of cash pressure visibly. Many small firms still lean on backward-looking summaries such as monthly totals or account balances. Those numbers describe what already happened. They rarely give enough lead time to respond before stress takes hold. The system tested here focuses on forward signals. It flags patterns that tend to appear before liquidity tightens, which gives managers space to adjust approvals, inventory decisions, or credit exposure while options remain open. This direction mirrors a wider movement in financial risk modeling, where forward-looking probability structures outperform static cutoffs. Makridis et al. (2023) frame this shift as a move toward detecting distributional change rather than chasing perfect point forecasts, placing more weight on early warning than on numerical precision [12]. The current study arrives at a similar conclusion in a smaller, operational setting.

Another outcome stands out. Effective working capital optimization does not depend on exotic models when the stress signal itself is clean. The strength of the logistic regression model suggests that SME liquidity strain often traces back to simple mechanics, especially sudden swings in revenue volatility. That matters for adoption. Smaller firms rarely have the appetite or resources for complex deep learning stacks. A transparent model that performs well lowers the entry barrier. The real contribution of AI here sits in the discipline of turning predictions into structured decisions. Early warnings feed into explicit approval rules. Those rules help smooth cash fluctuations, limit exposure during fragile periods, and protect resilience without falling into crude cost-cutting. The value comes from the loop between signal and action, not from model complexity for its own sake.

5.2. Customer experience trade-offs

One of the clearest lessons from the study comes from separating stress prediction from customer-level decisions. That separation makes it easier to see how customer experience survives even when liquidity tightens. The data show that customer behavior variables add little to the detection of firm-wide stress. They add a great deal once the firm decides how to respond. This split challenges a common instinct in modeling, where every feature is judged by its effect on predictive accuracy alone. Here, customer signals act more like steering controls. They shape intervention after a stress alert appears, allowing the firm to treat customers differently based on value and risk. The hybrid policy brings this idea to life. It avoids blanket restrictions by protecting access for customers who combine high value with low behavioral risk. That selectivity reduces avoidable revenue loss and softens the relational damage that follows indiscriminate denials. Similar patterns appear in supplier and credit risk research, where segmented, explainable decisions support long-term trust. Hasan et al. (2025) show that differentiated counterparty treatment strengthens supply networks by pairing caution with continuity [8]. In an SME setting, the same logic translates into loyalty that survives difficult weeks. Optimizing customer experience does not mean approving everything at all times. It means aligning restrictions with strategic customer value when resources are tightened.

5.3. Practical deployment considerations for U.S. SMEs

When you actually try to put this system to work, the first thing you notice is that it doesn't demand a big change to what SMEs are already doing. Most small businesses in the U.S. already have their sales and invoices tracked in some way, point-of-sale software, invoicing apps, or even a simple accounting program. The system only needs the basics: when a transaction happened, how much money came in, and which customer made it. From there, it quietly goes through the recent transactions, checks for signs of stress, and produces a score or probability. Those numbers then feed straight into a policy layer that lives inside the normal workflows. It can suggest approving or holding back certain transactions without forcing anyone to stop what they're doing or learn a whole new system. Another practical benefit is interpretability. Managers can see why stress alerts appear and understand how different policies influence revenue and customer approvals. That clarity makes it easier to trust the system and actually act on its recommendations. For smaller firms, where owners or managers are often personally involved in financial decisions, transparency is critical. Prior research shows that even technically accurate AI tools can face resistance if their logic feels opaque or hard to reconcile with intuition [3].

5.4. Limitations

It's important to be upfront about what this study does and does not capture. First, the stress labels are based on revenue volatility, not direct measures of cash shortfalls, insolvency, or missed obligations. They give an early warning signal but may not reflect every dimension of financial distress an SME could experience. Second, the analysis relies on a single retail dataset, which means the results might not generalize to other industries or business models. Service-oriented or manufacturing SMEs could show very different transactional patterns, so the effectiveness of the system could vary. Third, the dataset doesn't include detailed payment schedules, cost structures, or credit obligations, limiting the ability to model cash flow at a fine-grained level. These constraints are common in SME-focused AI research, where full financial visibility is rare, and proxies are often the only option [9]. Despite these gaps, the study demonstrates a clear proof of concept: it's possible to design AI-driven decision support that meaningfully balances financial stability with customer experience using data that's realistically available to small businesses.

6. Future work

Looking ahead, there are a few directions that could make this system more useful and closer to how SMEs actually operate. Right now, the framework uses revenue volatility as a stand-in for liquidity stress, but adding richer financial inputs could make a big difference. Things like payment delay distributions, cost structures, and short-term obligations would let the system see the difference between temporary dips in revenue and situations where cash positions are genuinely unsustainable. That kind of detail would improve both the predictions and the way policies are applied. You could think of it like adaptive urban design, where empty or transitional spaces are turned into functional, responsive areas through careful interventions. Abed (2025) talks about how context-aware design is key to making urban voids work [2]. SME financial systems could evolve similarly, responding dynamically to inflows, outflows, and timing mismatches rather than just looking at total revenue. Another important step is testing the framework beyond a single firm or industry. SMEs are far from uniform, and validating the model across different companies would help capture variations in operating cycles, customer bases, and risk tolerance. Trying it in different sectors like manufacturing, logistics, or professional services would show whether the approach holds up where cash flow patterns are very different. Multi-SME and cross-sector testing is essential for understanding which features matter most in each context. Shawon et al. (2025) emphasize that analytics built for resilience need to account for regional and sectoral differences to remain effective at scale [23]. Following that idea, future research should explore how stress definitions, volatility thresholds, and decision rules need to adapt across the SME landscape.

Looking beyond how well the system generalizes, there's a bigger question about cause and effect. Right now, we're basically judging the policies by running simulations. That tells us something, but it doesn't really separate the effect of the policy from patterns that were already in the data. Using tools like structural causal models or counterfactual evaluation could help tease that apart. It matters even more when small businesses are operating inside larger financial networks or using digital platforms and payment systems. Research in behavioral finance shows that what looks like a pattern in the data isn't always the real driver behind outcomes. Shawon and colleagues (2025) point out that spotting illicit activity in interconnected systems really requires modeling the causal pathways, not just chasing correlations [24]. If SME decision support could incorporate that kind of reasoning, it would be more trustworthy and less likely to produce unintended side effects. Then there's the question of explainability and designing for the people who actually use the system. The models we have now are interpretable, but that doesn't automatically mean a manager will understand what they're seeing. We could make outputs more intuitive with short narratives, scenario-based simulations, or interactive dashboards that let someone explore "what if" situations before deciding. Trust isn't just about the numbers being right; it's about being able to make sense of them in context. Future work should focus on matching explanations to how managers think, the culture of the organization, and any regulatory constraints. The goal is AI that actually helps people make better decisions, not something that tries to replace their judgment.

7. Conclusion

Working with SMEs, it quickly becomes clear that cash flow isn't just numbers on a spreadsheet; it's the lifeblood of daily operations. This study shows that AI can help make that lifeblood a little more predictable. By combining interpretable models like logistic regression with well-chosen financial and customer behavior signals, the system doesn't just spot problems; it helps managers act before stress turns into a crisis. It can suggest where to hold back, where to let things flow, and which customers matter most in keeping the business stable. What's interesting is that the hybrid approach, which weighs predicted stress against customer value, performs noticeably better than treating every transaction the same. It avoids repeated cash crunches without cutting off your most loyal clients. In practice, that's the difference between feeling reactive all the time and having a bit of breathing room to make decisions. There are caveats, of course. The stress signals here come from revenue fluctuations and inferred customer risk, so they don't cover everything; payment timing, costs, and some operational details aren't included. Plus, the study only looked at one SME dataset, so results might vary in another industry or region. That said, the approach is flexible. It can grow with more firms, more sectors, and richer data. It's not just a model on a screen; it's a framework that managers can use to tie predictions directly to decisions, which is where the real value lies.

The study contributes to accounting research by reframing liquidity risk as a predictive classification problem integrated with policy simulation. By aligning volatility forecasting with working capital optimization and customer retention dynamics, the framework bridges quantitative analytics and core accounting principles. This integration strengthens the theoretical foundation of AI-enabled financial management systems. At its core, this work isn't about proving AI can make perfect predictions. It's about showing that when predictions are connected to actionable rules, SMEs can manage cash flow more confidently, even under stress. It allows managers to make trade-offs thoughtfully, protecting revenue while keeping key customers engaged. There's plenty of room to take this further: bringing in more detailed financial inputs, testing across multiple businesses, and understanding the causal effects of policies could make these systems even more grounded in reality. But even as it stands, the approach offers a practical way for SME managers to move from reacting to anticipating, turning insights into real-world decisions.

References

[1] Aashish, K. C., Zamil, M. Z. H., Mridul, M. S. I., Akter, L., Sharmin, F., Ayon, E. H., … Malla10, S. (2025). Towards eco-friendly cybersecurity: Machine learning-based anomaly detection with carbon and energy metrics. International Journal of Applied Mathematics, 38(9s). https://doi.org/10.12732/ijam.v38i9s.822.

[2] Abed, J. (2025). Re-imagining the spaces in-between: Design interventions to transform urban voids in downtown Dayton into active, breathing areas. Journal of Mechanical, Civil and Industrial Engineering, 6(3), 84–97. https://doi.org/10.32996/jmcie.2025.6.3.10.

[3] Beg, M. O., Awan, M. N., & Ali, S. S. (2019). Algorithmic machine learning for the prediction of stock prices. In A. Rafay (Ed.), FinTech as a disruptive technology for financial institutions. IGI Global. https://doi.org/10.4018/978-1-5225-7805-5.ch007.

[4] Chouksey, A., Dola, A., Antara, U. K., Begum, S., Ahmed, T., Sultana, T., & Zabin, N. (2025). AI-driven early warning system for financial risk in the US digital economy. International Journal of Applied Mathematics, 38(9s). https://doi.org/10.12732/ijam.v38i9s.838.

[5] Das, B. C., et al. (2025). AI-driven cybersecurity threat detection: Building resilient defense systems using predictive analytics. arXiv preprint arXiv:2508.01422. https://doi.org/10.14419/hysdg957.

[6] Debnath, S., et al. (2025). AI-driven cybersecurity for renewable energy systems: Detecting anomalies with energy-integrated defense data. International Journal of Applied Mathematics, 38(5s). https://doi.org/10.12732/ijam.v38i5s.367.

[7] Enyiorji, P. (2025). Designing a self-optimizing cloud-native autonomous finance system for SMEs using multi-agent reinforcement learning. *International Journal of Financial Management and Economics*, *8*(1), 596-605. https://doi.org/10.33545/26179210.2025.v8.i1.660.

[8] Hasan, M. R., Rahman, M. A., Gomes, C. A. H., Nitu, F. N., Gomes, C. A., Islam, M. R., & Shawon, R. E. R. (2025). Building robust AI and machine learning models for supplier risk management: A data-driven strategy for enhancing supply chain resilience in the USA. Advances in Consumer Research, 2(4).

[9] Hasan, M. S., et al. (2025). Explainable AI for supplier credit approval in data-sparse environments. International Journal of Applied Mathematics, 38(5s). https://doi.org/10.12732/ijam.v38i5s.380.

[10] Jakir, T. (2025). Signal-to-noise analysis of crisis indicators in global finance using artificial intelligence. *International Journal of Applied Mathematics*, *38*(10s), 1815-1836. https://doi.org/10.12732/ijam.v38i10s.1075.

[11] Kotios, D., Makridis, G., Fatouros, G., & Kyriazis, D. (2022). Deep learning enhancing banking services: A hybrid transaction classification and cash flow prediction approach. Journal of Big Data, 9(1), 100. https://doi.org/10.1186/s40537-022-00651-x.

[12] Kotios, D., Makridis, G., Walser, S., Kyriazis, D., & Monferrino, V. (2022). Personalized finance management for SMEs. In Big Data and Artificial Intelligence in Digital Finance (pp. 215–232). Springer. https://doi.org/10.1007/978-3-030-94590-9_12.

[13] Makridis, G., Fatouros, G., Kotios, D., Soldatos, J., Filippakis, M., & Kyriazis, D. (2023). DeepVaR: A framework for portfolio risk assessment leveraging probabilistic deep neural networks. Digital Finance, 5(1), 29–56. https://doi.org/10.1007/s42521-022-00050-0.

[14] Malakauskas, A., & Lakštutienė, A. (2021). Financial distress prediction for small and medium enterprises using machine learning techniques. *Engineering Economics*, *32*(1), 4-14. https://doi.org/10.5755/j01.ee.32.1.27382.

[15] Medianovskyi, K., Malakauskas, A., Lakstutiene, A., & Yahia, S. B. (2022, July). Interpretable machine learning for SME financial distress prediction. In *International Conference on Computing and Information Technology* (pp. 454-464). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-031-25344-7_42.

[16] Mulier, K., Schoors, K., & Merlevede, B. (2016). Investment-cash flow sensitivity and financial constraints: Evidence from unquoted European SMEs. Journal of Banking & Finance, 73, 258–273. https://doi.org/10.1016/j.jbankfin.2016.09.002.

[17] Nguyena, H. A., & Nguyena, T. H. (2020). The prediction of future operating cash flows using accrual-based and cash-based accounting information: Empirical evidence from Vietnam. *Management Science Letters*, *10*, 683-694. https://doi.org/10.5267/j.msl.2019.9.010.

[18] Osuji, V. C., Dako, O. F., & Okafor, C. M. (2023). Seamless Integration of Digital Supply-Chain Platforms with Commercial Banking to Enhance Working Capital Efficiency for SMEs.

[19] Rafay, A. (Ed.). (2019). FinTech as a disruptive technology for financial institutions. IGI Global. https://doi.org/10.2139/ssrn.3376358.

[20] Rahman, M. S. (2025). Machine learning-enabled early warning system for detecting micro-inflation clusters in the US economy. International Journal of Applied Mathematics, 38(12s), 2743–2769. https://doi.org/10.12732/ijam.v38i12s.1585.

[21] Ray, R. K. (2025). Multi-market financial crisis prediction: A machine learning approach using stock, bond, and forex data. International Journal of Applied Mathematics, 38(8s), 706–738. https://doi.org/10.12732/ijam.v38i8s.602.

[22] Reza, S. A., et al. (2025). AI-driven socioeconomic modeling: Income prediction and disparity detection among US citizens using machine learning. Advances in Consumer Research, 2(4).

[23] Reza, S. A., Rahman, M. K., Rahman, M. D., Sharmin, S., Mithu, M. F. H., Hasnain, K. N., ... & Kabir, R. (2025). Machine Learning Enabled Early Warning System for Financial Distress Using Real-Time Digital Signals. *arXiv preprint arXiv:2510.22287*.

[24] Shawon, R. E. R., et al. (2025). Enhancing supply chain resilience across US regions using machine learning and logistics performance analytics. International Journal of Applied Mathematics, 38(4s). https://doi.org/10.12732/ijam.v38i4s.225.

[25] Shawon, R. E. R., et al. (2025). Detecting Illicit Cross-Chain Fund Movement: Behavioral Machine Learning Models for Bridge-Based Laundering Patterns. International Journal of Applied Mathematics, 38(12s). https://doi.org/10.12732/ijam.v38i12s.1399.
[26] Shovon, M. S. S. (2025). Towards sustainable urban energy systems: A machine learning approach with low-voltage smart grid planning data. International Journal of Applied Mathematics, 38(8s), 1115–1155. https://doi.org/10.12732/ijam.v38i8s.631.
[27] Sizan, M. M. H., et al. (2025). Machine learning-based unsupervised ensemble approach for detecting new money laundering typologies in transaction graphs. International Journal of Applied Mathematics, 38(2s). https://doi.org/10.12732/ijam.v38i2s.88.
[28] Xiao, J., Wen, Z., Jiang, X., Yu, L., & Wang, S. (2023). A three-stage research framework to assess and predict the financial risk of SMEs based on a hybrid method. Decision Support Systems. https://doi.org/10.1016/j.dss.2023.114090.
[29] Md Hasan Zamil. (2025). AI-driven business analytics for financial forecasting: A systematic review of decision support models in SMEs. Review of Applied Science and Technology, 4(02), 86-117. https://doi.org/10.63125/gjrpv442.