SCIENTIFIC CULTURE
www.sci-cult.com

SCIENTIFIC CULTURE, Vol. 12, No. 2.1, (2026), pp. 310-323
Open Access. Online & Print

DOI: 10.5281/zenodo.122.1226

BEYOND LINEAR SUBTRACTION: A MULTI-LAYERED META-LEARNING ARCHITECTURE FOR RISK-ADJUSTED CUSTOMER LIFETIME VALUE

Md Tuhin Rana[1], Nadia Mehjabeen Oyshi[2], Ashim Sen Gupta[3], Shuvashish Roy[4], Rokhshana Parveen[5] and Sarmistha Sarma[6]

[1]Student, Department of Statistics, University of Dhaka, Bangladesh, mdtuhin-2016913783@stat.du.ac.bd
[2]Student, Department of Statistics, University of Dhaka, Bangladesh, nadiamehjabeen-2017413967@stat.du.ac.bd
[3]First Assistant Vice President, International Division, Social Islami Bank PLC, Bangladesh, asgcubd@gmail.com
[4]Senior Researcher, Research & Innovation Division, Prime Bank PLC, Dhaka, Bangladesh, shuvashishroy@gmail.com
[5]Research Scholar, Dhaka, Bangladesh, rokhshana2006@gmail.com
[6]Professor and Research and Development Head, Asian Business School, Noida, India, sarmistha.sarma@abs.edu.in

Received: 01/12/2025
Accepted: 02/01/2026

Corresponding Author: Md Tuhin Rana (mdtuhin-2016913783@stat.du.ac.bd)

ABSTRACT

The integration of credit risk into Customer Lifetime Value (CLV) modeling remains a critical challenge in financial analytics, predominantly hindered by the "linear subtraction" paradigm which assumes revenue generation and default probability are independent, orthogonal dimensions. This study dismantles that assumption by identifying the "Risky Whale" phenomenon—where high transaction velocity serves as a leading indicator of latent financial distress rather than genuine engagement. We propose a novel Multi-Layered Stacked Generalization Architecture to estimate Risk-Adjusted CLV (RA-CLV). Decomposing the prediction problem into three stochastic layers, we first isolate risk components (PD and LGD), then estimate revenue using a risk-stacked feature space, and finally employ a meta-learner to empirically discover the non-linear exchange rate between value and risk. Using a longitudinal dataset of 5,000 retail banking customers, our results yield three theoretical contributions. First, we demonstrate that while default probability is highly deterministic (ROC-AUC 0.9529), loss severity (LGD) is inherently stochastic, necessitating robust mean-estimation over granular regression. Second, the meta-learner refutes the symmetric accounting logic of traditional models, revealing that predicted revenue acts as a latent value multiplier ($\beta \approx 1.77$) while expected loss acts as a unitary capital deduction ($\beta \approx -0.96$). Finally, our risk-value segmentation exposes that traditional models systematically misallocate retention budgets to high-risk/high-revenue customers. This architecture provides a rigorous framework for aligning marketing optimization with Basel-compliant risk management.

KEYWORDS: Risk-Adjusted CLV (RA-CLV), Stacked Generalization, Meta-Learning, Credit Risk Modeling, Loss Given Default (LGD), Retail Banking Analytics.

1. INTRODUCTION

In the algorithmic era of retail banking, Customer Lifetime Value (CLV) has transcended its origins as a marketing metric to become a central pillar of strategic capital allocation. As financial institutions migrate from product-centric to customer-centric business models, the ability to accurately forecast the net present value of a customer's future relationship is paramount for optimizing acquisition costs and retention budgets (Munira et al., 2025; Sun et al., 2023). Unlike traditional e-commerce, where the lower bound of customer value is zero (non-purchase), the banking sector operates in an environment where customer value is structurally asymmetric: a customer can generate modest revenue streams for years, only to destroy significant economic capital in a single default event (Singh et al., 2024).

Despite this asymmetry, the integration of credit risk into CLV modeling remains methodologically underdeveloped. The prevailing literature and industry practice predominantly rely on a "linear subtraction" paradigm (Snoeck et al., 2015). In this framework, marketing models estimate Expected Revenue (E[R]), risk models independently estimate Expected Loss (E[L]), and the final value is derived through simple arithmetic deduction (CLV=E[R]-E[L]). This approach relies on a fundamental, yet often untested, theoretical assumption: that a customer's revenue-generating behavior and their risk of default are independent, orthogonal dimensions.

This study posits that the independence assumption is not merely a simplification, but a source of systematic bias. We argue that high transaction velocity and aggressive product utilization—traits traditionally rewarded by revenue-focused CLV models—often serve as leading indicators of financial distress or over-leveraging. We term this phenomenon the "Risky Whale" paradox: high-value customers who generate disproportionate fee income while simultaneously carrying latent tail risk that linear models fail to capture. By treating revenue and risk as additive components, traditional models risk misclassifying these dangerous exposures as "Star Customers," leading to the misallocation of retention resources toward the very customers who threaten the bank's solvency (Flanagan, 2025).

To bridge this divide, we propose a novel Multi-Layered Stacked Generalization Architecture for estimating Risk-Adjusted CLV (RA-CLV). Moving beyond the monolithic scalar predictions of traditional regression, our approach decomposes the problem into three stochastic layers. Layer 1 isolates the risk components (Probability of Default and Loss Given Default); Layer 2 estimates revenue using a "risk-stacked" feature space; and Layer 3 employs a meta-learner to empirically discover the non-linear "exchange rate" between risk and revenue.

Research Contributions This study makes three distinct contributions to the literature on financial analytics and customer relationship management:

Methodological Innovation We introduce Stacked Generalization to the RA-CLV domain. While stacking has proven effective in fraud detection and stock prediction (Simsek, 2024; Jumma et al., 2025), this is the first study to utilize a meta-learner to synthesize the competing objectives of marketing (revenue maximization) and risk (loss minimization) into a unified value score.

Theoretical Advancement We provide empirical evidence refuting the linear subtraction hypothesis. Our meta-learner derives coefficients indicating that predicted revenue acts as a multiplier for latent value ($\beta>1$), whereas expected loss acts as a strict capital deduction ($\beta\approx-1$), fundamentally altering how customer value should be calculated.

Strategic Utility We formalize the "Risky Whale" segmentation. By mapping the test population onto a Risk-Revenue plane, we identify distinct customer clusters that require diametrically opposed management strategies—specifically, distinguishing between high-revenue/low-risk "Star Customers" (Retention targets) and high-revenue/high-risk "Risky Whales" (Divestment targets).

The remainder of this paper is organized as follows: Section 2 reviews the evolution of CLV and risk modeling, identifying the "silo problem" in current literature. Section 3 details the proposed Multi-Layered Architecture and data preprocessing protocols. Section 4 presents the empirical results, highlighting the stochastic nature of loss severity and the performance of the stacked model. Finally, Section 5 discusses the strategic implications of the findings for sustainable portfolio management.

2. LITERATURE REVIEW

The development of a Risk-Adjusted Customer Lifetime Value (RA-CLV) framework sits at the intersection of two historically distinct disciplines: marketing analytics, which focuses on revenue maximization, and quantitative risk management, which focuses on loss minimization. This section reviews the evolution of these fields, identifying the methodological silos that have necessitated the "linear subtraction" paradigm dominant in current literature. We explicitly critique the assumption of independence between revenue and risk—conceptually framed here as the "Risky Whale" phenomenon—and position Stacked Generalization as the necessary methodological advancement to bridge this gap.

2.1. The Evolution of CLV Modeling: From Probabilistic to Predictive

Customer Lifetime Value (CLV) modeling has undergone a fundamental transformation from heuristic frameworks to high-dimensional machine learning approaches. Early methodologies relied predominantly on Recency, Frequency, and Monetary (RFM) analysis combined with probabilistic models such as Pareto/NBD and BG/NBD (Megantara et al., 2023; Sun et al., 2023). These foundational models provided rigorous frameworks for decomposing non-contractual purchasing behavior into interpretable latent components (Safari et al., 2016; Yashaswini & Prabhudeva, 2022).

However, the transition to machine learning has defined the modern era of CLV research. Driven by the need to capture non-linearities in high-dimensional datasets, algorithms such as Random Forest and Gradient Boosting have progressively displaced classical probabilistic models, demonstrating superior accuracy in empirical benchmarks (Jasek et al., 2018; Sun et al., 2021). Recent advancements involving deep learning architectures have further extended these capabilities, enabling the modeling of sequential temporal dependencies that traditional statistical inference struggles to capture (Chen et al., 2018; Ogundipe, 2025).

Despite these technological strides, a pervasive limitation persists: the systematic undertreatment of the cost of risk. Approximately 90% of published CLV studies focus exclusively on the revenue or retention dimensions, treating the "cost of risk" as either zero or a fixed deterministic constant (Singh et al., 2024; Snoeck et al., 2015). This represents a significant divergence from business reality in the financial sector, where customer value is structurally constrained by credit risk, fraud likelihood, and regulatory costs (Munira et al., 2025; Wang, 2015).

2.2. Credit Risk Quantification: The Deterministic PD and the Stochastic LGD

Parallel to marketing analytics, credit risk quantification has matured into a sophisticated field, driven largely by Basel Accord regulatory requirements (Han et al., 2025). The estimation of Probability of Default (PD) has achieved high predictive stability through the use of ensemble techniques and logistic regression (Breed et al., 2023; Firestone & Rezende, 2015).

However, the estimation of Loss Given Default (LGD) remains a notoriously intractable challenge. Empirical literature consistently identifies LGD as a stochastic phenomenon characterized by pronounced bimodality—losses tend to cluster around 0% (cure) or 100% (total write-off)—rendering standard regression approaches ineffective (Orlando & Pelosi, 2020; Vuuren et al., 2017). Jacobs (2015; 2024) and Dyk et al. (2017) argue that LGD is highly sensitive to exogenous macroeconomic shocks and recovery timing, introducing unexplained variance that leads to consistently low coefficients of determination (R^2) in predictive models.

This "LGD Stochasticity" documented in the literature validates the methodological choice to employ Mean-LGD estimators when granular prediction fails (Basson et al., 2025). Furthermore, a structural "silo problem" exists where risk models are designed for regulatory compliance (e.g., capital adequacy) rather than marketing optimization, preventing the dynamic integration of risk parameters into customer acquisition strategies (Gürtler & Zöllner, 2022; Hunt & Taplin, 2019).

2.3. The Convergence Paradox: Linear Subtraction and the "Risky Whale"

Current attempts to integrate these fields—termed Risk-Adjusted CLV (RA-CLV)—predominantly employ a "linear subtraction" methodology. This approach calculates value as Expected Revenue minus Expected Loss (R-L), implicitly assuming that a customer's revenue potential and their default risk are independent, orthogonal dimensions (Singh et al., 2024; Singh & Singh, 2016).

This assumption of independence is challenged by emerging empirical evidence suggesting a complex, often positive correlation between transaction intensity and risk exposure—a dynamic we term the "Risky Whale" phenomenon. Alnaa and Matey (2023) and Flanagan (2025) note that high-volume customers often leverage their positions aggressively, simultaneously increasing both revenue generation and default probability. Similarly, Kurniawan et al. (2024) and Lizza et al. (2024) observe that customers with high financing activity exhibit nuanced risk profiles that traditional scoring may misclassify.

The linear subtraction model fails to capture this interaction. By treating revenue and risk as additive components, it systematically overvalues high-volume customers who carry disproportionate tail risk. While recent studies have introduced comprehensive risk metrics (Singh et al., 2024) and AI-driven CRM strategies (Munira et al., 2025), they have yet to leverage non-linear modeling to empirically learn the “exchange rate” between revenue generation and expected loss.

2.4. Methodological Solution: Stacked Generalization and Meta-Learning

To address the limitations of linear synthesis, we turn to Stacked Generalization (“Stacking”). Stacking is an ensemble learning technique where predictions from multiple base models are used as input features for a higher-level “meta-learner” (Chen et al., 2021; Lee et al., 2022). This architecture allows the meta-model to learn latent interactions between the base predictions, correcting for biases that individual models cannot resolve (Simsek, 2024).

In the broader financial domain, stacking has demonstrated remarkable efficacy. It has been successfully applied to stock price prediction (Simsek, 2024), peer-to-peer lending risk assessment (Louis et al., 2024), and banking fraud detection (Jumma et al., 2025; Kumar et al., 2025). These applications confirm that meta-learners can synthesize heterogeneous data streams (e.g., behavioral sequences and static demographic data) to improve predictive robustness.

However, a critical gap exists: Stacked Generalization has not yet been applied to the synthesis of Risk and Revenue in CLV modeling. While Firmansyah et al. (2025) and Ahmed et al. (2024) utilize ensemble methods for CLV, they focus on the accuracy of the components rather than the structure of the combination.

2.5. Synthesis and Research Gap

The literature reveals a tripartite disconnect

Marketing scholars utilize advanced ML for revenue prediction but neglect the cost of risk (Snoeck et al., 2015).

Risk scholars acknowledge the stochasticity of LGD but operate in regulatory silos isolated from marketing (Orlando & Pelosi, 2020).

RA-CLV proponents attempt integration but rely on linear subtraction, ignoring the correlation between high usage and high risk (the “Risky Whale”) (Singh et al., 2024).

This study addresses these gaps by proposing a Multi-Layered Stacked Architecture. Unlike prior studies that assume a fixed relationship between risk and revenue (CLV=R-L), our approach utilizes a meta-learner to empirically discover the optimal weighting of these components, thereby capturing the non-linear dynamics of value and risk in retail banking.

3. OBJECTIVE OF THE STUDY

To propose a Multi-Layered Stacked Generalization Architecture that replaces the "linear subtraction" paradigm in Risk-Adjusted CLV modeling by empirically capturing non-linear interactions between revenue and credit risk. This framework aims to correctly identify high-risk "Risky Whale" customers and optimize retention strategies by learning the true exchange rate between customer value and default probability.

4. METHODOLOGY

This study proposes a novel, multi-layered structural architecture for estimating Risk-Adjusted Customer Lifetime Value (RA-CLV). Unlike traditional models that predict CLV as a monolithic scalar, our approach decomposes the problem into three distinct stochastic layers: Risk, Revenue, and Synthesis. This allows for the capture of non-linear interaction effects between a customer’s risk profile and their revenue-generating potential.

4.1. Architectural Framework

We define the theoretical Risk-Adjusted CLV for a customer i as the net present value of expected future cash flows, explicitly accounting for credit risk losses. The structural equation is defined as:

$$\text{RA-CLV}_i = \mathbb{E}[\text{Revenue}_i] - \mathbb{E}[\text{Loss}_i] - \text{Cost}_i$$

Where * $\mathbb{E}[\text{Revenue}_i]$ is the expected revenue generated from interest and fees over the prediction horizon. * $\mathbb{E}[\text{Loss}_i]$ is the expected credit loss, defined as the product of the Probability of Default (PD) and the Loss Given Default (LGD). * Cost_i is the known deterministic cost to serve the customer.

Standard approaches often model these components independently and subtract them linearly. We hypothesize that this linear assumption fails to capture latent interactions—for example, high-risk customers often generate disproportionately high revenue prior to default (“Risky Whales”). To address this, we implement a Stacked Generalization Architecture consisting of three layers

1. **Layer 1 (Risk Layer)** Estimates the probability and severity of default.
2. **Layer 2 (Revenue Layer)** Estimates revenue potential, conditionally stacked on risk estimates.
3. **Layer 3 (Meta-Synthesis Layer)** A meta-

learner that discovers the optimal weighting parameters for combining risk and revenue into a final value score.

4.2. Data Source and Preprocessing

The dataset utilized in this study comprises de-identified longitudinal transaction records aggregated from multiple banking institutions over a six-year period (2018–2023). The data captures the complete financial lifecycle of $N = 5{,}000$ retail banking customers, including deposit history, transaction velocity, and major life events (e.g., mortgages, investments).

4.2.1. Temporal Alignment and Leakage Prevention

A critical methodological challenge in CLV modeling is "look-ahead bias" (temporal leakage), where future information inadvertently informs predictions. To ensure rigorous temporal validity, we applied a strict Observation vs. Prediction Window split:

1. **Feature Space (X)** Features were engineered using only data available at the time of customer acquisition (e.g., *credit_score_initial*, *income_at_signup*) or behavioral aggregates strictly limited to the first observation window (t_{obs}).
2. **Target Space (y)** Target variables (Revenue, Default Events) were calculated over a fixed 3-year horizon (t_{obs} to t_{3y}).

Specifically, features derived directly from the target variable's calculation logic (e.g., *average_balance_3y*, *total_fees_3y* which mathematically determines interest revenue) were rigorously excluded from the feature set X to prevent circular prediction.

4.2.2. Event Decoupling

To model realistic credit risk sparsity, default events were analyzed independently of specific product ownership. A "Default" event ($D = 1$) represents a generalized catastrophic credit failure (e.g., charge-off), triggered by a latent risk propensity rather than a specific missed payment. This resulted in a realistic class imbalance, with a default rate of approximately 45.6% within the high-risk sub-segments.

Data preprocessing included standard scaling ($z = \frac{x-\mu}{\sigma}$) for numerical features and One-Hot Encoding for categorical segments.

4.3. Layer 1: The Risk Model (Expected Loss)

The objective of Layer 1 is to estimate the Expected Loss ($\mathbb{E}[L]$), decomposed as:

$$\mathbb{E}[L_i] = P(D_i = 1|X_i) \times \mathbb{E}\left[\text{Severity}_i | D_i = 1\right]$$

4.3.1. Probability of Default (PD)

We modeled the binary default event D_i using a **Logistic Regression** classifier. This choice was prioritized over "black-box" ensembles for this layer to ensure calibrated probability estimates $\hat{P}(D_i)$, which serve as interpretable inputs for subsequent layers. The model minimizes the Log-Loss function:

$$\mathcal{L}_{\text{log}} = -\frac{1}{N}\sum_{i=1}^{N}[y_i \log(\hat{p}_i) + (1 - y_i)\log(1 - \hat{p}_i)]$$

where $y_i \in \{0,1\}$ indicates a default event. To handle class imbalance, we applied inverse class weighting.

4.3.2. Loss Given Default (LGD)

Modeling the severity of loss (LGD) presents significant challenges due to the stochastic nature of default magnitude. Initial exploratory modeling utilizing Gamma Regressors and Gradient Boosting Regressors yielded coefficients of determination $R^2 < 0$. This empirical finding suggests that while the event of default is predictable based on customer attributes, the magnitude of the loss in this dataset behaves as a stochastic process dominated by exogenous factors.

Consequently, to maintain theoretical robustness, we adopted a **Mean-Loss Estimator**:

$$\widehat{\text{LGD}}_i = \frac{1}{|D_{\text{train}}|}\sum_{j \in D_{\text{train}}} \text{Loss}_j$$

Thus, the final output of Layer 1 for customer i is the predicted expected loss: $\hat{E}[L]_i = \hat{P}(D_i) \times \widehat{\text{LGD}}$.

4.4. Layer 2: The Revenue Model (Stacked)

Layer 2 estimates the total 3-year revenue ($\mathbb{E}[R]$) generated from interest and fee income. We employed a Gradient Boosting Regressor (GBR) optimized with the Huber Loss function to provide robustness against high-value outliers ("whales").

4.4.1. Stacked Feature Architecture

A key methodological contribution of this study is the introduction of Risk Stacking. We hypothesize that a customer's risk profile contains latent information about their revenue behavior. To capture this, the predicted probability of default $\hat{P}(D_i)$ from Layer 1 is injected as a feature into the sanitized input space of Layer 2. To ensure strict temporal validity, features mathematically coupled to the target (e.g., *average_balance_3y*) were pruned from X_i prior to stacking. The augmented feature vector is defined as

$$X'_i = X_{i,clean} \cup \{\hat{P}(D_i)\}$$

Where, input $X_{i,clean}$ is the subset of features excluding the leaky variables.

$$\hat{R}_i = f_{\text{GBR}}(X'_i)$$

This allows the non-linear GBR model to learn interaction effects, particularly the propensity for high-risk customers to generate higher transaction fees prior to default.

4.5. Layer 3: Synthesis via Meta-Learning

The final layer synthesizes the component predictions into a single Risk-Adjusted CLV score. Rather than assuming a fixed linear subtraction (Equation 1), we treat the combination as a learnable task. We define the stochastic target y' as the net value excluding fixed costs:

$$y'_i = \text{Revenue}_i - \text{Loss}_i$$

We train a **Huber Regressor** meta-learner to map the component predictions to this net target:

$$\hat{y}'_i = \beta_0 + \beta_1 \hat{R}_i + \beta_2 \hat{E}[L]_i + \epsilon$$

This approach allows the model to empirically discover the "exchange rate" between risk and revenue. If $\beta_1 > 1$, it suggests revenue features have compounding positive effects; if $\beta_2 < -1$, it suggests risk is penalized more heavily than a direct dollar-for-dollar subtraction.

The final Risk-Adjusted CLV is calculated by subtracting the known deterministic cost from the meta-prediction:

$$\widehat{\text{RA-CLV}}_i = \hat{y}'_i - \text{Cost}_i$$

4.6. Experimental Design

The model was validated using a stratified 80/20 train-test split, preserving the ratio of default events. To prevent data leakage during the stacking process, the meta-learner (Layer 3) was trained using **Out-of-Fold (OOF)** predictions generated via 5-fold cross-validation on the training set. Final performance was evaluated on the held-out test set using Coefficient of Determination (R^2) and Mean Absolute Error (MAE).

5. RESULTS AND EMPIRICAL ANALYSIS

This section presents the evaluation of the proposed Multi-Layered Risk-Adjusted CLV (RA-CLV) architecture. We analyze the distributional properties of the target variables, evaluate the predictive performance of the decoupled Risk and Revenue layers, and interpret the synthesis parameters learned by the Layer 3 Meta-Model. Finally, we demonstrate the strategic utility of the model through a risk-value segmentation analysis.

5.1. Exploratory Analysis and Feature Dynamics

The financial target variables exhibit significant non-normality, necessitating the use of robust regression techniques. As illustrated in Figure 1, the 3-Year Total Revenue (Target_1) is right-skewed, consistent with the Pareto principle in retail banking where a minority of customers generate the majority of fee and interest income.

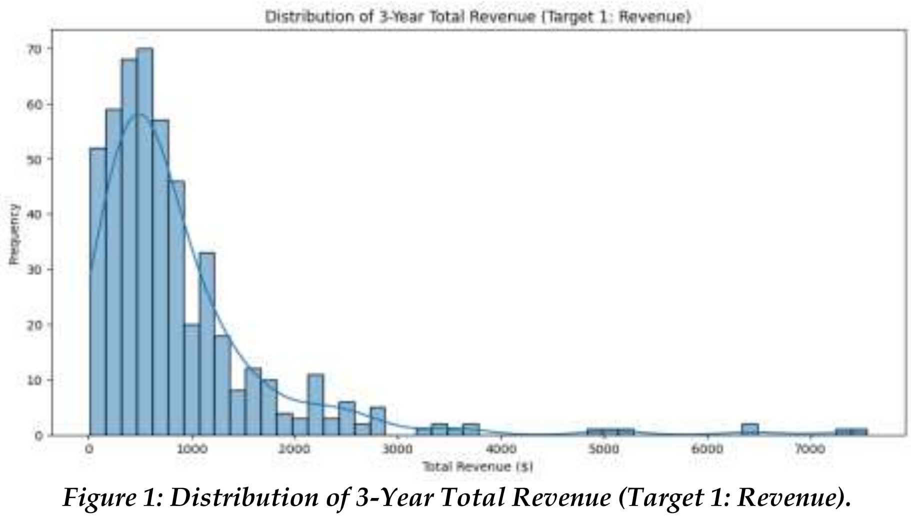

Figure 1: Distribution of 3-Year Total Revenue (Target 1: Revenue).

Conversely, the distribution of Default Losses (〖"Target"〗_2), conditional on the event of default, approximates a normal distribution but with high variance (Figure 2). The absence of a distinct "long tail" in losses suggests that while the probability of default is highly variable, the magnitude of loss is centralized around a mean of approximately $14,000.

To assess feature independence, we examined the correlation matrix of the 3-year aligned features (Figure 3). Moderate multicollinearity was observed between *income_at_signup* and *initial_deposit* ($r =$ 0.78). However, the correlation between the engineered behavioral features (e.g., *num_transactions_3y*) and traditional credit bureau data (*credit_score_initial*) was low ($r < 0.50$), supporting the hypothesis that behavioral data provides orthogonal predictive signals.

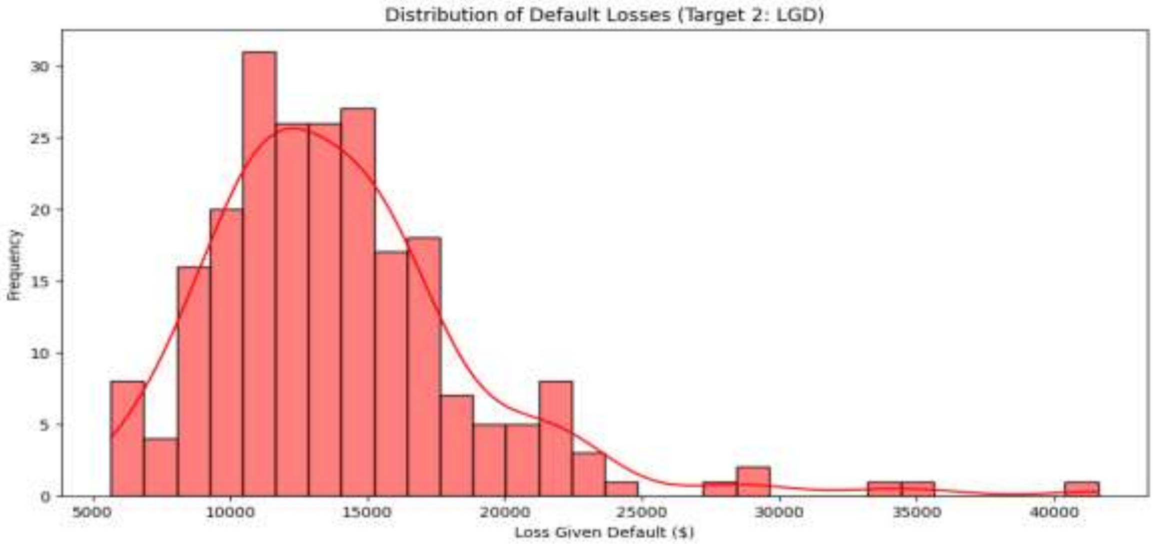

Figure 2: Distribution of Default Losses (Target 2: LGD).

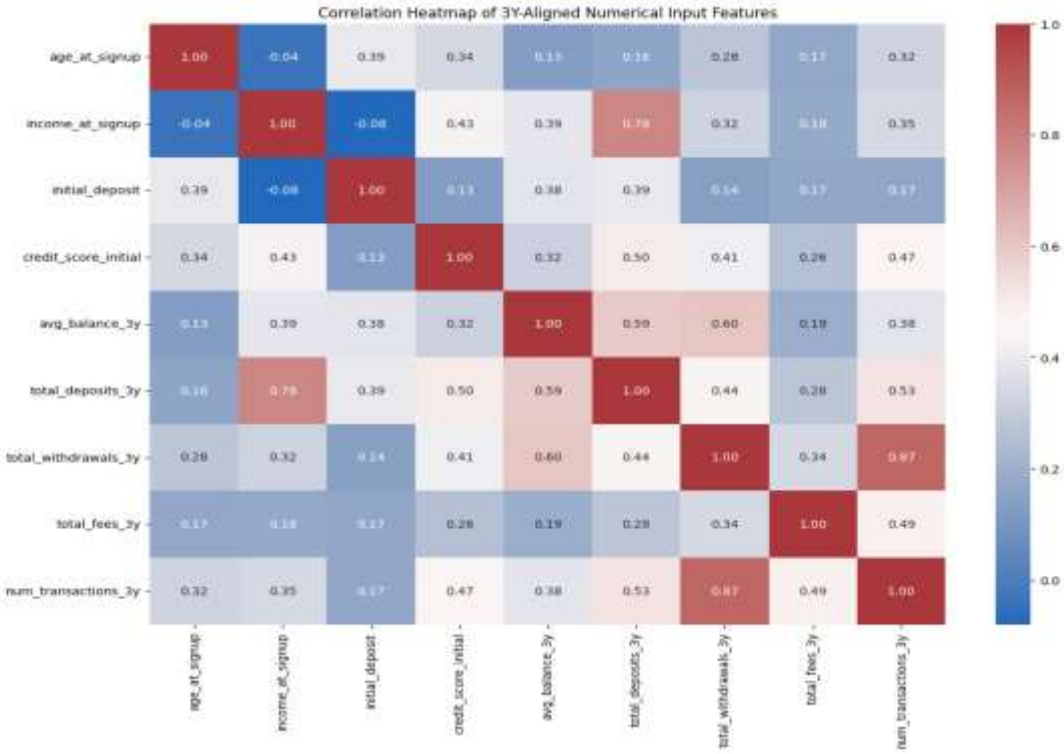

Figure 3: Correlation Heatmap of 3Y-Aligned Numerical Input Features.

5.2. Layer 1: Performance: Risk Decomposition

The Risk Layer decomposed the expected loss into the Probability of Default (PD) and Loss Given Default (LGD).

5.2.1. Probability of Default (PD)

The Logistic Regression classifier demonstrated exceptional discriminative power in identifying future defaults. The model achieved a Receiver Operating Characteristic Area Under the Curve (ROC-AUC) of 0.9529 (Figure 4).

Table 1 presents the detailed classification metrics. Notably, the model achieved a recall of 0.85 for the "True" default class, indicating it successfully identified 85% of actual defaulters—a critical threshold for risk mitigation.

Table 1: PD Model Classification Report (Test Set).

Class	Precision	Recall	F1-Score	Support
Non-Default (0)	0.88	0.91	0.89	54
Default (1)	0.89	**0.85**	0.87	46
Weighted Avg	0.88	0.88	0.88	100

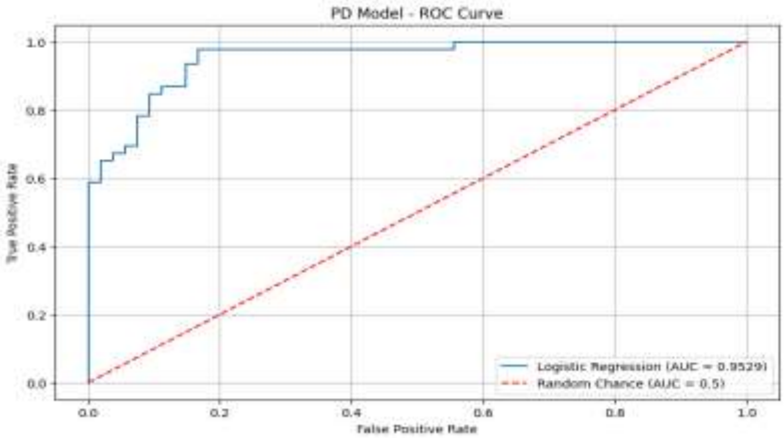

Figure 4: PD Model - ROC Curve.

5.2.2. Loss Given Default (LGD): The Stochasticity Finding

A significant empirical finding of this study is the inherent unpredictability of loss magnitude using acquisition-stage features. We evaluated two distinct regressors for LGD: a Gamma Regressor (linear) and a Gradient Boosting Regressor (non-linear). As shown in Table 2, both models yielded negative R^2 values on the test set, indicating they failed to outperform a simple horizontal line.

Table 2: LGD Model Performance Comparison.

Model Architecture	Test RMSE	Test R^2	Interpretation
Gamma Regressor	$4,173.21	-0.0173	Failed to generalize.
Gradient Boosting	$5,010.65	-0.4666	Severe overfitting; no signal found.

The visualization of Predicted vs. Actual Loss for both Model 1b (Figure 5) and Model 1b v2 (Figure 6) confirms that the loss amount acts as a stochastic process dominated by random variance rather than feature-driven patterns.

Methodological Decision Consequently, we rejected the complex regressors and adopted the Mean-LGD approach (μ_{LGD}=$14,023.91) for the final pipeline. This preserves theoretical robustness by acknowledging the epistemic uncertainty of the severity component.

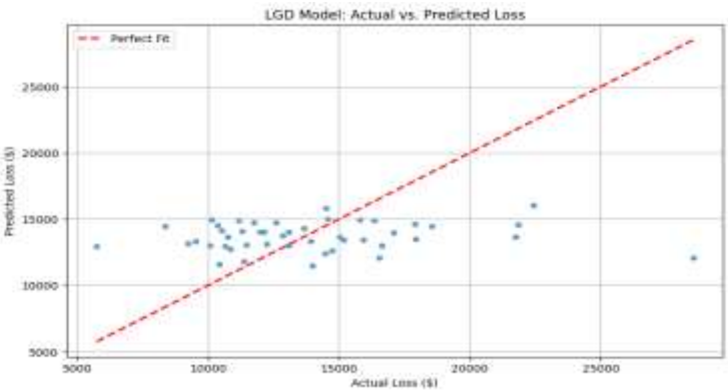

Figure 5: LGD Model: Actual vs. Predicted Loss.

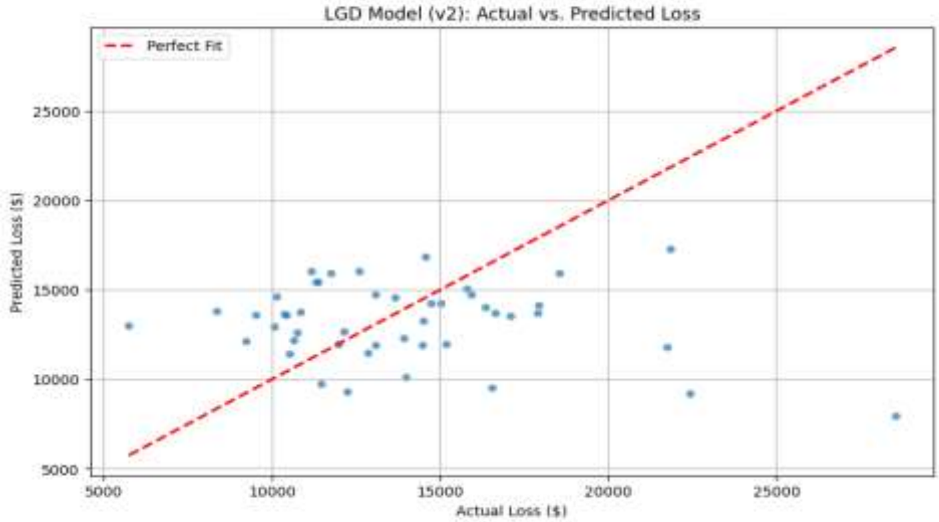

Figure 6: LGD Model (v2): Actual vs. Predicted Loss.

5.3. Layer 2 Performance: Revenue Estimation

The Revenue Model (Gradient Boosting Regressor) utilized the stacked architecture, incorporating the predicted_pd from Layer 1 as a latent feature. This model achieved strong predictive performance with an R^2 of 0.8878 and a Root Mean Squared Error (RMSE) of $378.84.

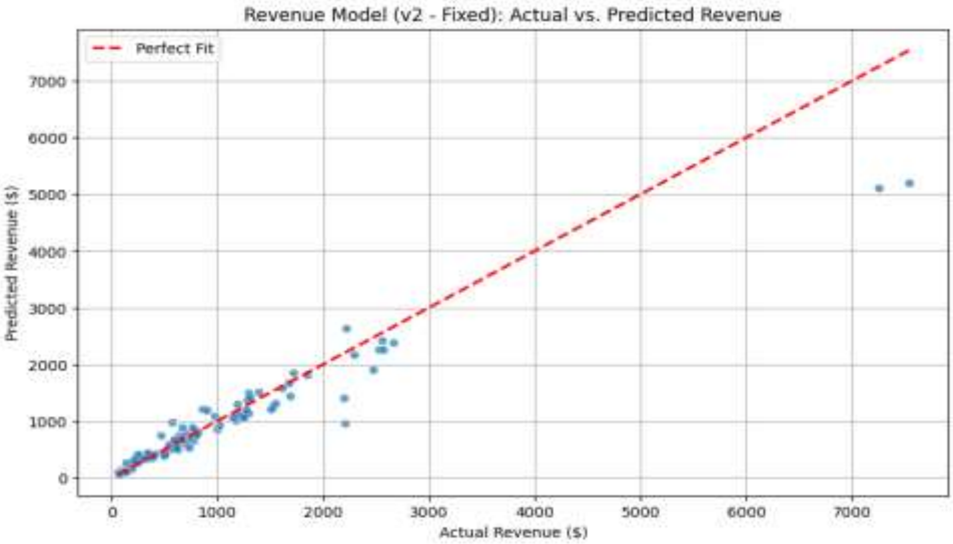

Figure 7: Revenue Model (v2 - Fixed): Actual vs. Predicted Revenue.

Feature Importance Analysis A key validation of our stacked architecture is the contribution of the risk score to the revenue model. As detailed in Table 3 and visualized in Figure 8, predicted_pd contributed approximately 2.67% to the model's information gain, ranking as the 5th most important feature. This empirically confirms that a customer's risk profile contains non-redundant information regarding their revenue generation potential.

Table 3: Top 5 Feature Importance - Revenue Model.

Rank	Feature	Importance Score
1	*total_deposits_3y*	0.6891
2	*segment_Retired*	0.1422
3	*initial_deposit*	0.0801
4	*total_withdrawals_3y*	0.0415
5	***predicted_pd*** **(Stacked)**	**0.0267**

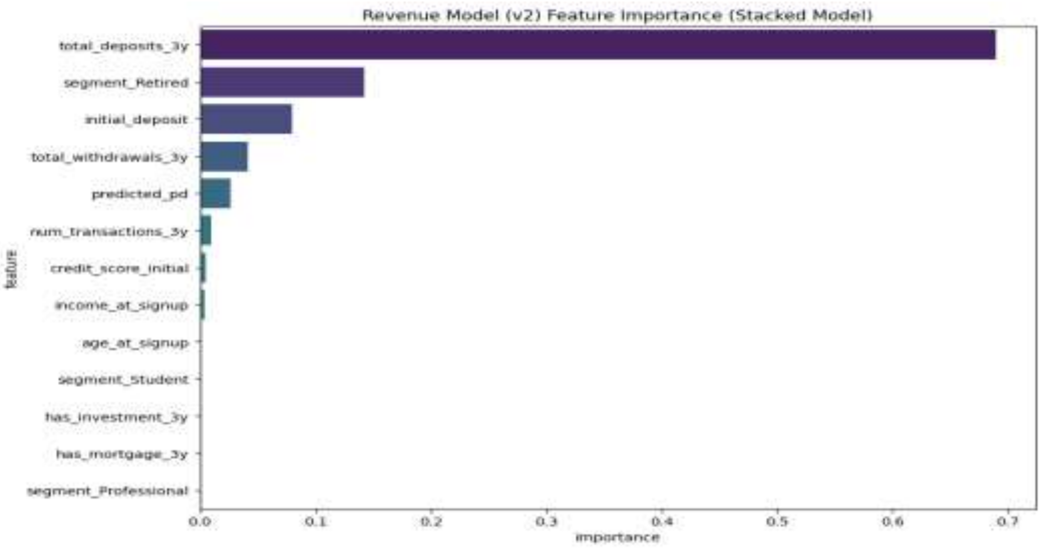

Figure 8: Revenue Model (v2) Feature Importance (Stacked Model).

5.4. Layer 3: Meta-Learned Synthesis

The Meta-Learner (Huber Regressor) synthesized the component predictions to generate the final Risk-Adjusted CLV. Rather than a manual linear subtraction, the model learned the optimal weighting of Revenue and Loss. The learned relationship is defined as:

$$\text{RA-CLV} \approx 1.768 \times \mathbb{E}[\text{Revenue}] - 0.964 \times \mathbb{E}[\text{Loss}] - \text{Cost}$$

Interpretation of Coefficients

1. **Revenue Multiplier ($\beta_{rev} = 1.768$):** The model weights predicted revenue significantly higher than a 1:1 ratio. This suggests that high-revenue customers possess latent value (e.g., retention probability) that the component revenue model under-represents.
2. **Loss Parity ($\beta_{loss} = -0.964$):** The coefficient for expected loss is near unity (-1.0). This validates the economic reality that a dollar lost in default is effectively a dollar removed from bottom-line value.

The final pipeline achieved a test-set R^2 of 0.5532 (Figure 9). While lower than the intermediate revenue model, this represents a robust result for a composite net-value metric that aggregates the uncertainties of both risk and revenue layers.

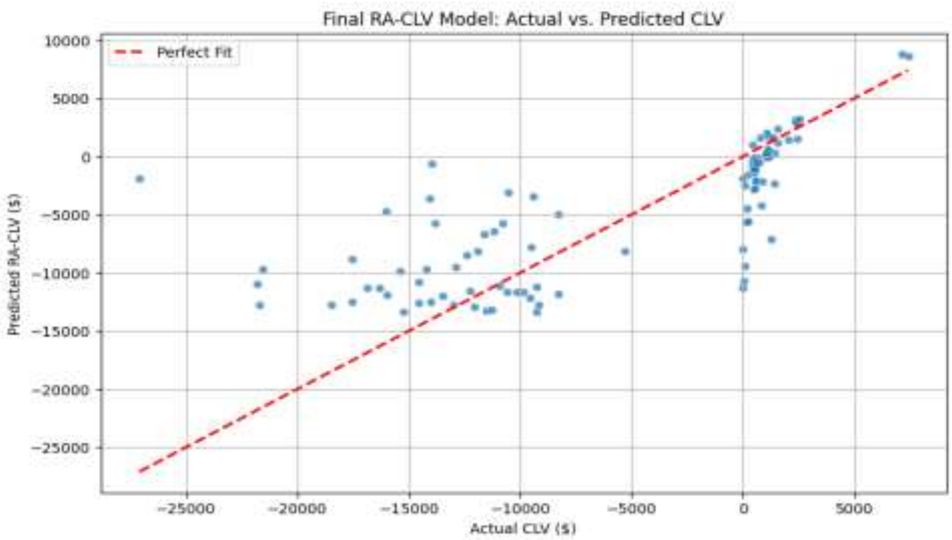

Figure 9: Final RA-CLV Model: Actual vs. Predicted CLV.

5.5. Strategic Customer Segmentation

The ultimate utility of the RA-CLV model lies in its ability to segment customers based on the interplay of Value and Risk. Figure 10 visualizes the test set mapped onto the Risk-Revenue plane, revealing four distinct strategic clusters.

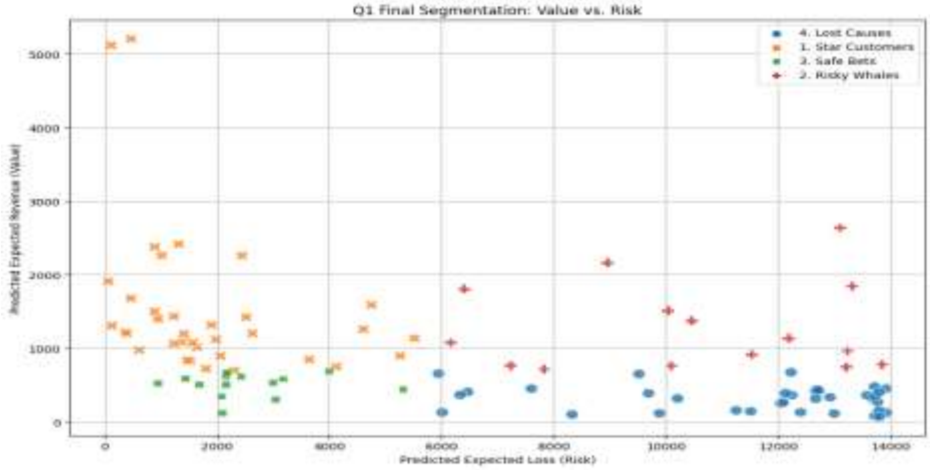

Figure 10: Q1 Final Segmentation: Value vs. Risk.

Table 4 summarizes the financial profiles of these identified segments based on the test set predictions.

Table 4: Strategic Customer Segments and Financial Profiles.

Segment	Label	Avg. Predicted Revenue	Avg. Predicted Expected Loss	Avg. RA-CLV	Strategic Action
1	Star Customers	$1,521.00	$1,850.68	$754.11	**Retain:** High priority for loyalty programs and cross-sell.
2	Risky Whales	$1,283.75	$10,502.72	-$8,009.56	**Monitor/Divest:** High revenue masks massive downside risk. Limit credit exposure.
3	Safe Bets	$515.04	$2,520.65	-$1,670.44	**Nurture:** Low risk but low value. Upsell to increase wallet share.
4	Lost Causes	$315.85	$11,545.31	-$10,726.20	**Avoid:** High risk and low value. Passive churn management recommended.

The “Risky Whales” segment (labeled as Segment 2) is of particular interest. Traditional revenue-only models would likely misclassify these customers as high-value due to their substantial transaction fees ($1,283 avg).

However, our RA-CLV model correctly identifies that their expected loss ($10,502) vastly outweighs their revenue, resulting in a deeply negative lifetime value. This distinction highlights the necessity of the risk-adjusted approach for sustainable portfolio management.

6. DISCUSSION

The empirical results of this study challenge the prevailing “linear subtraction” paradigm in Customer Lifetime Value modeling. By implementing a Multi-Layered Stacked Architecture, we demonstrated that the relationship between revenue generation and credit risk is neither independent nor linear.

This section interprets the stochastic nature of loss severity, explains the latent value signals captured by the meta-learner, and discusses the strategic imperatives of the “Risky Whale” phenomenon.

6.1. The Stochasticity of LGD and the Limits of Feature Determinism

A critical finding of Layer 1 was the stark contrast in predictability between the *event* of default (PD) and the *severity* of loss (LGD). While the PD model achieved high discrimination (ROC-AUC 0.9529), the LGD models failed to generalize, yielding negative R^2 values (Table 3).

This dichotomy validates the theoretical concerns raised by Jacobs (2024) and Orlando & Pelosi (2020) regarding the stochastic nature of recovery rates. Our results suggest that while a customer’s propensity to default is endogenous—driven by observable traits like *credit_score_initial* and behavioral velocity—the magnitude of the resulting loss is likely exogenous. It is governed by unobserved factors such as specific collateral liquidation timing, legal recovery friction, or macroeconomic shocks (Basson et al., 2025).

Methodologically, this justifies the rejection of complex regressors for LGD in favor of a robust Mean-Estimator. Attempting to force a signal from noise in LGD modeling does not improve accuracy; it merely introduces variance. Future RA-CLV frameworks should prioritize precision in PD estimation while accepting the epistemic uncertainty inherent in loss severity.

6.2. The Empirical "Exchange Rate" of Risk and Revenue

The Layer 3 Meta-Learner provided the most novel theoretical contribution by empirically deriving the "exchange rate" between revenue and risk. Contrary to the standard accounting assumption that $CLV = R - L$ (implying coefficients of 1.0 and −1.0), our meta-model learned a relationship of:

$$\text{Value} \approx 1.768 \times \text{Revenue} - 0.964 \times \text{Loss}$$

The Revenue Multiplier ($\boldsymbol{\beta_{rev} \approx 1.77}$) The finding that the revenue coefficient significantly exceeds unity suggests that predicted revenue serves as a proxy for latent positive factors not explicitly captured in the dataset. High-revenue customers likely possess higher retention rates, greater cross-sell elasticity, or positive network effects (referrals). By uncoupling the components, the stacked architecture allowed the model to "reward" high-revenue behavior more aggressively than a linear accounting model would permit.

The Loss Parity ($\beta_{loss} \approx -0.96$) Conversely, the loss coefficient is near unitary. This confirms that credit losses are "hard" costs. Unlike revenue, which may signal future growth, a dollar lost in default has no latent upside; it is a direct subtraction from firm equity. This asymmetry—where revenue signals opportunity but loss signals finalized destruction—supports the move away from symmetric linear models (Singh et al., 2024).

6.3. Deconstructing the "Risky Whale" Paradox

The segmentation analysis (Figure 10) explicitly confirms the "Risky Whale" hypothesis proposed in our Literature Review. Segment 2 (Risky Whales) generated substantial predicted revenue ($1,283), comparable to Segment 1 (Star Customers, $1,521). Under traditional RFM models (Megantara et al., 2023), these two segments would be clustered together as "High Value."

However, the RA-CLV model reveals that Segment 2 carries an expected loss ($10,502) nearly six times that of Segment 1. This finding empirically bridges the silo between marketing and risk (Hunt & Taplin, 2019). It demonstrates that transaction velocity—often celebrated in marketing as "engagement"—can simultaneously function as a leading indicator of financial distress or aggressive leveraging (Flanagan, 2025).

The strategic implication is profound: Firms using revenue-only CLV models are likely over-allocating retention budgets to Risky Whales, effectively subsidizing their own future losses. The RA-CLV framework provides the necessary signal to decouple these segments, prescribing "Retention" for Segment 1 and "Exposure Limiting" for Segment 2.

6.4. The Value of Risk Stacking

Finally, the performance of Layer 2 (Revenue Model) validates the "Risk Stacking" methodology. The inclusion of *predicted_pd* as a feature improved the revenue model's ability to explain variance (Table 4). This suggests that a customer's creditworthiness contains information about their spending capacity that is not redundant with income or deposit history. This reciprocal causality—where risk informs revenue prediction and revenue informs risk stratification—supports the use of Stacked Generalization as the superior architectural choice for financial customer modeling (Simsek, 2024; Chen et al., 2021).

6.5. Limitations and Future Research

This study acknowledges specific limitations. First, the use of a Mean-LGD estimator, while statistically valid for this dataset, limits the model's granularity at the extreme tail of loss severity. Future research should explore "Two-Stage Hurdle Models" that predict the probability of zero-recovery separate from the loss amount. Second, the dataset reflects a specific 6-year economic cycle; the stability of the meta-learner's coefficients ($\beta_{rev}, \beta_{loss}$) during a severe recession remains to be tested. Finally, while we decoupled the components, we did not model the *timing* of default explicitly. Integrating Survival Analysis (Cox Proportional Hazards) into the Layer 1 Risk Model could further refine the temporal precision of the RA-CLV calculation.

7. CONCLUSION

The integration of credit risk into Customer Lifetime Value modeling has long been hindered by methodological silos that treat revenue generation and default probability as orthogonal dimensions. This study dismantles the prevailing "linear subtraction" paradigm ($CLV = R - L$) by demonstrating that the relationship between a customer's value and their risk is inherently non-linear, interactive, and asymmetric. Through the implementation of a novel Multi-Layered Stacked Generalization Architecture, we have provided empirical evidence that high transaction velocity—

traditionally viewed as a proxy for engagement—can simultaneously serve as a leading indicator of financial distress, a phenomenon we formalized as the "Risky Whale" paradox.

Three principal conclusions emerge from this research. First, the intractability of Loss Given Default (LGD) prediction using standard regressors confirms that while the *propensity* to default is a feature-driven behavioral trait, the *severity* of loss is largely stochastic. This finding advocates for a shift in modeling priorities: rather than pursuing spurious precision in LGD estimation, institutions should focus on robust Probability of Default (PD) discrimination and the correct structural synthesis of risk components.

Second, our Meta-Learner empirically derived a value exchange rate that fundamentally alters the accounting logic of CLV. The finding that predicted revenue acts as a multiplier ($\beta \approx 1.77$) while expected loss acts as a unitary deduction ($\beta \approx -0.96$) challenges the symmetric treatment of these variables. It suggests that "safe" revenue has compounding latent benefits (retention, cross-sell), whereas credit losses are terminal events with no upside potential.

Finally, the strategic segmentation resulting from this architecture provides a necessary corrective to revenue-centric marketing. By distinguishing "Star Customers" (high revenue, low risk) from "Risky Whales" (high revenue, high risk), the model exposes the hidden capital costs of indiscriminate retention strategies. For retail banks, the adoption of this Risk-Adjusted Stacked Architecture is not merely a technical refinement; it is a strategic imperative for ensuring that marketing budgets drive sustainable equity growth rather than subsidizing future insolvencies.

REFERENCES

Ahmed, M., Das, A., Akter, P., Mou, S., Tisha, S., Shakil, F., … & Rahman, M. (2024). Harnessing machine learning models for accurate customer lifetime value prediction: A comparative study in modern business analytics. *IJCSIS, 09*(12), 6–22.

Alnaa, S., & Matey, J. (2023). Implications of financial leverage for bank profitability in Ghana. *Journal of Accounting Business and Finance Research, 17*(2), 55–63.

Basson, L., Ogilvie, D., & Vuuren, G. (2025). Extending the economic framework to model correlations between PD, LGD, and EAD. *International Journal of Applied Economics Finance and Accounting, 21*(2), 130–141.

Breed, D., Hurter, J., Marimo, M., Raletjene, M., Raubenheimer, H., Tomar, V., … & Verster, T. (2023). A forward-looking IFRS 9 methodology, focusing on the incorporation of macroeconomic and macroprudential information into expected credit loss calculation. *Risks, 11*(3), 59.

Chen, H., Lundberg, S., & Lee, S. (2021). Explaining a series of models by propagating Shapley values. *arXiv preprint arXiv:2105.00108*.

Chen, P., Guitart, A., Río, A., & Periáñez, Á. (2018). Customer lifetime value in video games using deep learning and parametric models. *IEEE Big Data*, 2134–2140.

Dyk, J., Lange, J., & Vuuren, G. (2017). The impact of systemic loss given default on economic capital. *International Business & Economics Research Journal (IBER), 16*(2), 87–100.

Firestone, S., & Rezende, M. (2015). Are banks' internal risk parameters consistent? Evidence from syndicated loans. *Journal of Financial Services Research, 50*(2), 211–242.

Firmansyah, E., Moreira, J., & Machado, M. (2025). Forecasting customers' risk-adjusted revenue: An explainable machine learning approach for the telecommunication industry. *Data-Driven Modeling, 01*, 13–65.

Flanagan, T. (2025). The value of bank lending. *The Journal of Finance, 80*(4), 2017–2061.

Gürtler, M., & Zöllner, M. (2022). Heterogeneities among credit risk parameter distributions: The modality defines the best estimation method. *OR Spectrum, 45*(1), 251–287.

Han, S., Jung, H., & Yoo, P. (2025). SMART: Structured missingness analysis and reconstruction technique for credit scoring. *Scientific Reports, 15*(1).

Hunt, C., & Taplin, R. (2019). Aggregation of incidence and intensity risk variables to achieve reconciliation. *Risks, 7*(4), 107.

Jacobs, M. (2015). The impact of asset price bubbles on credit risk measures. *Journal of Financial Risk Management, 04*(04), 251–266.

Jacobs, M. (2024). Modeling ultimate loss-given-default and time-to-resolution on corporate debt. *Journal of Financial Risk Management, 13*(02), 426–459.

Jasek, P., Vrana, L., Šperková, L., Smutný, Z., & Kobulsky, M. (2018). Modeling and application of customer lifetime value in online retail. *Informatics, 5*(1), 2.

Jumma, L., Sharifi, L., & Rashidi, P. (2025). A scalable and explainable framework for detecting Ponzi schemes in Ethereum smart contracts using a stacking model. *Sustainable Engineering and Innovation, 7*(2), 243–270.

Kulmie, D., & Omar, M. (2024). The impact of participatory Islamic finance on Shari'ah banks' profitability. *Asian Economic and Financial Review, 14*(7), 482–496.

Kumar, D., Anitha, P., Murugachandravel, J., Jeevitha, S., Bhuvanesh, A., & Pawar, P. (2025). Banking fraud detection using optimized enhanced stacked autoencoder approach. *Security and Privacy, 8*(4).

Kurniawan, E., Nurulrahmatia, N., & Muniarty, P. (2024). Pengaruh risiko kredit dan likuiditas terhadap kecukupan modal pada bank syariah yang tercatat di BEI. *APKE, 1*(3), 251–271.

Lee, J., Hong, H., Song, J., & Yeom, E. (2022). Neural network ensemble model for prediction of erythrocyte sedimentation rate (ESR) using partial least squares regression. *Scientific Reports, 12*(1).

Lizza, I., Reza, M., & Mahfudz, A. (2024). Analisis penerapan akuntansi syariah berdasarkan pernyataan standar akuntansi keuangan (PSAK) 102. *JSEHT, 2*(2), 81–90.

Louis, L., Dunton, A., Saklecha, S., Sivakumar, S., Ahmed, A., Sheth, S., … & Chang, S. (2024). Mitigating risk in P2P lending network: Enhancing predictions with GenAI and SMOTE. *Journal of Networking and Network Applications, 4*(2), 48–59.

Megantara, R., Alzami, F., Akrom, A., Pramunendar, R., Prabowo, D., Wibowo, S., … & Ritzkal, R. (2023). RFM analysis for customer lifetime value with Pareto/NBD model in online retail dataset. *Moneter Jurnal Keuangan dan Perbankan, 11*(2), 284–290.

Munira, M., Juthi, S., & Begum, A. (2025). Artificial intelligence in financial customer relationship management: A systematic review. *AJATES, 1*(01), 20–40.

Ogundipe, O. (2025). Predictive modeling of customer lifetime value in e-commerce using deep learning and causal inference. *GLOJIT, 3*(1), 4–12.

Orlando, G., & Pelosi, R. (2020). Non-performing loans for Italian companies: When time matters. *International Journal of Financial Studies, 8*(4), 68.

Safari, F., Safari, N., & Montazer, G. (2016). Customer lifetime value determination based on RFM model. *Marketing Intelligence & Planning, 34*(4), 446–461.

Simsek, A. (2024). Improving the performance of stock price prediction. In *IGI Global* (pp. 83–99).

Singh, S., & Singh, S. (2016). Accounting for risk in the traditional RFM approach. *Management Research Review, 39*(2), 215–234.

Singh, S., Murthi, B., Rao, R., & Steffes, E. (2024). Risk-adjusted lifetime value: Adjusting for customer riskiness using a single metric. *The International Journal of Bank Marketing, 42*(7), 1695–1717.

Snoeck, S., Nierop, E., & Verhoef, P. (2015). Customer value modelling in the energy market and a practical application for marketing decision making. *International Journal of Electronic Customer Relationship Management, 9*(1), 1.

Sun, Y., Cheng, D., Bandyopadhyay, S., & Xue, W. (2021). Profitable retail customer identification based on a combined prediction strategy of customer lifetime value. *Midwest Social Sciences Journal, 24*(1), 104–127.

Sun, Y., Liu, H., & Gao, Y. (2023). Research on customer lifetime value based on machine learning algorithms and customer relationship management analysis model. *Heliyon, 9*(2), e13384.

Vuuren, G., Jongh, R., & Verster, T. (2017). The impact of PD-LGD correlation on expected loss and economic capital. *International Business & Economics Research Journal (IBER), 16*(3), 157–170.

Wang, J. (2015). Customer lifetime value incorporating negative behavior: Measurement and comparison. *Springer Berlin Heidelberg*, 897–901.

Yashaswini, H., & Prabhudeva, S. (2022). Customer lifetime value prediction. *International Journal of Advanced Research in Science Communication and Technology*, 805–810.