Review Article *Open Access*

Journal of Economics and Social Dynamics

ISSN: 3069-6577

DOI: *doi.org/10.63721/26JESD0148*

CreditSentinel AI: An ML-Augmented Credit Risk Framework for Emerging Markets

William Fernandes

Department of Mechatronics, SVKM's NMIMS Mukesh Patel School of Technology Management and Engineering, India

Citation: *William Fernandes (2026) CreditSentinel AI: An ML-Augmented Credit Risk Framework for Emerging Markets. J of Eco and Soc Dynamics. 2(4), 1-7. WMJ/JESD-148*

Abstract

Traditionally, the process of credit risk adjudication has relied on periodic financial disclosures, supplemented with heuristic-based analyses. Although the process offers some level of oversight, the approach has significant drawbacks, such as information latency and cognitive bias. These issues are further compounded by the nature of the markets, which are known to be volatile, with opaque financial disclosures and significant governance risks. In the present research, the authors introduce CreditSentinel AI, a high-quality, machine learning-based approach to solvency analysis. The proposed CreditSentinel AI model utilizes non-linear eXtreme Gradient Boosting (XGBoost), incorporating high-frequency market data and fundamental financial statement-based financial information. The key differentiator of the CreditSentinel AI approach is the Graduated Penalty Mechanism (GPM), which utilizes severity-weighted risk adjustments, as opposed to binary covenant-based risk adjustments. In the proposed research, the CreditSentinel AI approach was tested using stress tests on listed firms on the National Stock Exchange (NSE), demonstrating the proposed approach's capability to detect structurally distressed capital profiles and leverage-based default risks. CreditSentinel AI has the potential to become a high-performance, automated first-line risk sentinel, capable of reducing decision latency.

***Corresponding author**: William Fernandes, Department of Mechatronics, SVKM's NMIMS Mukesh Patel School of Technology Management and Engineering, India.

Submitted: 02.04.2026 **Accepted**: 06.04.2026 **Published**: 25.04.2026

Keywords: Credit Risk Modeling, Machine Learning, XGBoost, Emerging Markets, Corporate Governance, Fintech

Introduction

In today's credit markets, the rate at which information is assimilated can be as critical as the accuracy of the analytical judgment. Institutional lenders and corporate credit desks make investment decisions on the basis of prospective solvency indicators rather than relying on past historical performance. However, the overwhelming majority of existing credit underwriting models are grounded in periodic disclosures and rely on manual data extraction and spreadsheet-based ratio analysis. Credit analysts generally await periodic disclosures such as quarterly (equivalent to 10-Q/10-K) and annual disclosures and manually recast the financials before assessing the risk profile of the issuer.

The antiquated methodology adopted in assessing companies' financial health gives rise to a critical latency problem. Corporate distress is a rare event and is generally triggered by market events such as a fall in the stock price of the firm, liquidity issues, margin calls, and refinancing risk. By the time such risks are reflected in the financial disclosures of the companies, creditors generally end up holding distressed debt rather than possessing warning indicators. This is particularly true in the case of emerging markets (EM), wherein disclosure lags, accounting diversity, and governance opacity further decrease the effectiveness of traditional models.

Moreover, existing credit risk models, such as discriminant score models [1] and structural distance to default models [2], were largely developed for developed markets with welldefined ownership structures. However, when implemented in emerging markets like India, existing traditional models tend to underestimate the risks and ignore regional-specific governance risks. One of the most significant examples of the inability of traditional models to account for the realities of emerging markets can be demonstrated through the sharepledging phenomenon of promoters to raise personal loans. This results in a vicious cycle between the stock price and the creditors of the firm, wherein a fall in stock price causes the creditors to demand payment of loans, thereby causing the firm to sell shares.

The CreditSentinel AI model has been developed to overcome the deficiencies of traditional credit risk models. CreditSentinel AI has been developed to replicate the process of institutionalized credit decision-making through the use of non-linear machine learning.

The contributions of this paper are threefold:

- Real-Time Adjudication: A pipeline that bridges live market data with fundamental accounting metrics.
- Non-Linear Modeling: The application of XGBoost to capture cliff-edge risks that linear models tend to miss.
- Governance Integration: The quantification of Promoter Pledging as a primary solvency variable.

Theoretical Framework and Related Work

The discipline of automated credit risk modeling has, until recently, been dominated by two main theoretical approaches: the structural approach, based on the principles of option pricing theory, and the statistical approach, based on the principles of discriminant analysis. Although these approaches have shown significant promise for modeling credit risk in the context of the developed markets, with their associated high levels of liquidity and transparency, they exhibit some significant shortcomings when applied to the context of the high-beta, governance-driven emerging markets such as India.

Structural Models: The Limitations of Merton (1974) Equity, according to the credit risk model of Robert C. Merton, can be considered a call option on the assets of the firm, where the strike price of the option is the face value of the firm's debt (D).In theory, default occurs if the market value of a its assets (V_A) at maturity(T) falls below its outstanding debt.
The Mechanic: The model calculates the Distance to Default (DD) by estimating the volatility of the firm's assets (σ_A).

$$DD = \ln(V_A/D) + (\mu - 0.5\sigma_A^2)T / (\sigma_A\sqrt{T})$$

The Merton model for credit risk assessment faces problems when applied to emerging markets (EM), driven by its underlying efficient market hypothesis, which includes the assumptions of perfect reflection of asset values by equity prices and the use of volatility as a measure of risk. These assumptions are violated in EM markets by three related factors:

- Illiquidity: The mid-cap equities of emerging markets often suffer from low float, where a large sale may dominate price movements rather than insolvency.
- Noise vs. signal: The high volatility of EM equities may be driven by macroeconomic sentiment, such as currency fluctuations, rather than credit risk.
- Unobservable volatility: The estimation of asset volatility (σ_A) requires inference from equity volatility, which may be artificially kept low in markets with significant manipulation and pledges, only to increase suddenly and explosively

when a credit event occurs, a 'jump diffusion' process, which may cause the Merton model to underestimate credit risk until a credit event occurs.

CreditSentinel AI approach: The CreditSentinel AI model avoids the problem of market-implied asset values, which are often volatile, and instead uses Cash Flow Serviceability (EBITDA). This is a more robust measure than asset volatility, which is noisy, particularly in illiquid markets, making CreditSentinel AI more robust for credit stressed issuers.

Statistical Models: The Linearity Problem of Altman Z-Score (1968)

Edward Altman's Z-Score remains the industry standard for bankruptcy prediction. It employs Linear Discriminant Analysis (LDA) to separate distressed firms from healthy ones using a weighted linear combination of five ratios:

$$Z = 1.2X_1 + 1.4X_2 + 3.3X_3 + 0.6X_4 + 1.0X_5$$

The variable X_3 (EBIT/Total Assets) often appears as the key driver.

The Linearity Trap describes the major weakness in the Zscore model as perceived in institutional underwriting.

As a way to describe the weakness, the following scenarios should be considered:

- A firm that increases its leverage from 2.0x to 3.0x has a fixed numerical change in its Z-score model.
- A firm that increases its leverage from 5.0x to 6.0x has the same numerical change in its Z-score model.

Justification for Incorrectness

The non-linear nature of credit risk follows an exponential curve as one approaches a default scenario. A firm increasing its leverage from 2x to 3x often has a chemically neutral effect, while a firm increasing its leverage from 5x to 6x often triggers cross-default provisions, rating agency downgrade, and liquidity issues. The model does not encapsulate the "Cliff Effect" as a linear model would, instead re-classifying a terminal entity as merely "moderately risky."

CreditSentinel AI Methodology

The use of XGBoost (Extreme Gradient Boosting) replaces the traditional linear weights with decision trees. The nonlinear nature of decision trees inherently determines split points (IF Debt > 5.0x THEN Risk High), which can encapsulate the non-linear escalation of risk that occurs at specific covenant triggers.

The Governance Void: The "Alpha" of Promoter Pledging

Merton's market-based model and Altman's accountingbased model are unable to take adequate account of Agency Risk, specifically the conflict between the interests of the controlling shareholder (Promoters) and the debt holders. The key driver of this issue in Asian markets involves promoters using personal holdings of shares as collateral to raise capital for unrelated business ventures. This creates a reflexive process that creates a unique risk transmission mechanism not found in the traditional models used in the West.

Sequence of Process

- Pledge of shares worth 40% of the total shares as collateral.
- A decline in the share price of the firm by 15% due to market noise.
- A call for more collateral from the lender to sell shares.
- Sale of shares, which absorbs a large volume of shares in the market, leading to a decline in the share price of the firm by a further 20%.
- Difficulty in raising capital for the firm; lenders also tighten credit conditions.

The blind spot here is that a firm may have an ideal ZScore (profitable with low corporate debt) but still be on the verge of collapse due to the high level of promoter leverage. Research conducted by the Reserve Bank of India (RBI) has specifically highlighted high levels of pledging as a key driver of default risk, although this factor is still not reflected in solvency models.

CreditSentinel AI

The above model addresses this theoretical blind spot by making the Promoter Pledging Percentage a key input variable (X_6). This model can now predict "Death Spirals" that are not reflected in financial models.

Methodology: Data & Feature Engineering

The key concept of CreditSentinel AI is to replicate the cognitive process of an experienced senior credit officer, simultaneously addressing the statistical challenges inherent in the process of prognostication. This is done through a feature engineering process that converts simple accounting information into risk-informed signals to overcome the challenges of small data.

Feature Engineering: The Four Pillars of Solvency
This selection process aligns with the conceptual model of the "5 Cs of Credit," which include Character, Capacity, Capital, Collateral, and Conditions. However, this process is also quantitatively and dimensionally limited. While traditional data sets have hundreds of ratios to draw from, CreditSentinel AI limits itself to only four dimensions. This limits overfitting and maintains interpretability, which is a key feature that is critical to ensuring Model Risk Management compliance.

Leverage: Net Debt / EBITDA (Capacity)
This metric is the core covenant in corporate credit agreements and institutional lending. This metric gives a temporal solvency test, which estimates the number of years it takes to pay off net debt based on current operating cash flows, assuming zero growth and zero capital expenditures.

Reasoning: Unlike using Debt/Equity (subject to capitalization rate swings) or Debt/Assets (subject to bookvalue accounting), Net Debt/EBITDA focuses solely on operating repayment capability. This metric is the least susceptible to manipulation through accounting leverage.

Equation
Net Debt / EBITDA = (Total Long-Term Debt + Short-
Term Debt – Cash & Equivalents) / EBITDA

Threshold Dynamics: In general, for industrial companies, a debt to EBITDA ratio of around 4.0x is utilized, and for utilities and infrastructure, this figure increases to around 6.0x. These figures represent the baseline for the penalty function in CreditSentinel AI.

Serviceability: Interest Coverage Ratio (Cash Flow)
While leverage measures debt levels, serviceability measures the impact of debt levels. The interest coverage ratio measures the margin between earnings and necessary interest payments.

Rationale: This metric is critical in determining what are referred to as "Zombie Companies," i.e., companies with earnings levels to cover interest payments but insufficient earnings to cover principal payments. In this case, a coverage ratio of less than 1.5x is universally seen as a negative indicator, regardless of debt levels.

Equation: Interest Coverage Ratio (ICR) = Earnings Before Interest and Taxes (EBIT) / Interest Expense

Liquidity: Current Ratio (Short-Term Resilience)
While liquidity concerns have a shorter-term focus, solvency issues have a long-term focus. The Current Ratio was created to answer a simple question: Will the firm survive for the next twelve months without requiring refinancing?

Rationale: In emerging markets, credit availability can suddenly disappear (e.g., the IL&FS crisis in India). The firm may have good long-term solvency fundamentals (e.g., a low Debt/EBITDA ratio) but may have poor short-term liquidity fundamentals, with impending commercial paper maturities and poor cash balances. The Current Ratio is a "Veto Variable" that can override other good fundamentals if it dips below 1.0x.

Equation: Current Ratio = Current Assets / Current Liabilities

Governance: Promoter Pledging Percentage (Character/Agency Risk)
This metric was created to capture a unique attribute of the system and measure the "Character" risk of the firm. In India, promoters of companies often pledge their shares to increase their leverage.

Rationale: High levels of pledged shares represent a nonoperating channel for default risk. When the share price dips to a certain level, the promoters' pledged shares can be invoked and sold, thereby driving down the share price and triggering a "death spiral" of margin calls. This risk is outside the balance sheet and is of a catastrophic nature.

System Architecture

The CreditSentinel AI framework is designed and deployed as a modular cloud-native application running atop the Google Colab runtime environment. This allows for better reproducibility by offering instant access to preconfigured Python-based Scientific Computing Stacks (Pandas, XGBoost) and cloud based computation for execution of the code.

This architecture follows a linear data processing pipeline based on the Extract-Transform-Load-Analyze (ETLA) pattern and is divided into three different layers:

Layer 1: Fault-Tolerant Data Harvester

At the center of the CreditSentinel AI is the Data Harvester. It is a robust data ingestion tool designed to work with the Yahoo Finance API using the yfinance library. The key challenge associated with modeling emerging market equities is the sparsity of financial data. For mid-cap companies, financial data is not always fully tagged using the eXtensible Business Reporting Language (XBRL) or timely with quarterly reports. As a solution, the CreditSentinel AI framework utilizes a novel approach called the Permissive Substitution Logic (PSL) integrated into the Fault-Tolerant Data Harvester.

Ticker Normalization and Sanitization

Before taking the financial data from the Yahoo Finance API using the yfinance library, the user inputs undergo normalization and sanitization.

- Suffix Handling: The CreditSentinel AI framework automatically appends the '.NS' suffix for the National Stock Exchange (NSE) tickers if not provided by the user.
- Currency Standardization: Financial information is standardized to Indian Rupees (INR) to avoid currency mismatch issues during the analysis of dual-listed entities.

The Waterfall Substitution Algorithm The conventional process may break down if there are any occurrences of Null values in any critical field. CreditSentinel AI uses a hierarchical substitute for the conventional process.

This includes the following steps:

- Step 1 (Primary Query): CreditSentinel AI retrieves the latest quarterly (Q-0) equivalent financial statements.
- Step 2 (Null Detection): A logic gate detects any occurrence of Null values in any critical solvency metric: Total Debt, EBITDA, or Interest Expense.
- Step 3 (Secondary Fallback): If any defects are detected in the data, CreditSentinel AI calculates the Trailing Twelve Months (TTM) aggregate by adding the results from the last four quarters.
- Step 4 (Tertiary Fallback): If the TTM calculation fails, CreditSentinel AI uses the latest Annual Report.

Benefits

The hierarchical process results in a continuity rate of 99.9%. This process differentiates between a "broken company" and "broken data." This prevents the model from discarding a valid entity due to missing information.

Layer 2: The Logic Engine (The "Brain") Once the cleaning stage is complete, the dataset is processed and fed into the Logic Engine. The Logic Engine uses a dual-track processing method, where probabilistic machine learning inference and deterministic rule-based computation occur simultaneously.

Track A: The Probabilistic Classifier (XGBoost)

Vectorization: The raw financial information is converted to a feature vector x, which is defined as x = [L, ICR, CR, QoQ].

Inference: The feature vector x is utilized to calculate a prediction using a pre-trained Extreme Gradient Boosting Classifier, referred to hereafter as XGBoost. The main advantage of using XGBoost over other linear regressionbased methods lies in the Gradient Boosting method, which uses an ensemble of decision trees to accommodate nonlinear relationships, referred to as "cliff effects." For instance, if the current ratio changes from 1.2 to 1.0, this change is more significant than if it changes from 2.2 to 2.0.

Output: The output of this process is a raw probabilistic score $P\_distress \in [0,1]$, representing the statistical probability of the firm going through financial distress.

Track B: The Deterministic Penalty Engine (GPM)

The “Hard Constraints” of the institutional risk policy are imposed through the application of the “Graduated Penalty Mechanism” (GPM) rules.

Sector-Specific Logic: This engine has dynamic thresholds depending on the sector. More “headroom” is allowed in the “Utilities” sector compared to the “Manufacturing” sector.

Severity Calculation: This engine calculates the severity of the breach of the covenant (Δ). A 5% breach results in a “Penalty” for “Watchlist,” and a breach greater than 50% results in a “Penalty” for “Distressed.”

Aggregation

The final risk score is calculated as a combination of the results from the two tracks.

Risk_final = min(99%, P_distress + Σ Penalties)

It makes use of the capabilities of Artificial Intelligence and maintains the security features of traditional credit policy at the same time.

Layer 3: Visualization & Interaction (The Frontend) User interface: It has been created using the Streamlit library, a Python-based tool for the development of reactive web applications.

Methodological approach: In terms of information density and velocity of decision-making, the methodological approach adopted for the development of the dashboard aligns perfectly with the characteristics of a Bloomberg Terminal.

Traffic Light Risk Gauge

Complex probabilistic calculations have been simplified for the end user through a simple and user-friendly radial gauge, avoiding information overload for the user.

- 0-30%: Investment Grade / Safe
- 30-70%: Speculative / Monitor
- 70-100%: Distressed / Avoid

This configuration enables the analyst to quickly and effectively ascertain the information provided.

Velocity and Trend Analysis

A static ratio does not provide any useful information, and the velocity of the change in the ratio is more relevant. For the development of the dashboard, the Plotly.js library has been utilized for the development of an interactive chart for the “Leverage Velocity,” representing the Net Debt/EBITDA for the latest four periods. This methodological approach enables analysts to differentiate between fallen angels and distressed issuers.

Empirical Evaluation

The model was also subjected to evaluation using a realworld set of National Stock Exchange (NSE) listed firms. This validates the model’s capability to identify discriminating signals in real-world scenarios.

Safety Validation: ABB India Profile: Conservative Capital Goods Producer.

Observed Metrics:

- Net Debt/EBITDA: 0.04x (Negligible)
- Interest Coverage: 114.15x (Exceptional)
- Pledging: 0.63% (Clean) Model Output:
- Base Risk: 1% (Investment Grade)
- Penalties: None Analysis:

The model correctly interprets a triple-digit metric. This metric is indicative of good health from a financial perspective. Moreover, the logarithmic normalization process also prevents the positive outlier effect from the 114x coverage ratio. Thus, this example supports the model’s reliability.

Distress Detection: Adani Green Energy Profile: Capital-intensive Infrastructure Developer.

Observed Metrics:

- Net Debt/EBITDA > 9.0x (Critical Breach)
- Sector Limit: 6.0x (Utilities)
- Serviceability: Weak Coverage Model Output:
- Base Risk: High (Due to Raw Leverage)
- Leverage Penalty: "Leverage: Excessive (>25% Risk)" triggered due to Graduated Penalty Mechanism (Δ > 0.5)
- Interaction Penalty: "Structure: Compound Risk Detected
- (+10%)" triggered due to Graduated Penalty Mechanism (Δ

Ø 0.5)

- Final Score: 99% (Distressed/Default) Analysis:

The output complies with the CreditSentinel Institutional Credit covenant rules. While equity holders have the freedom to grow their business, creditors would not approve a 9x leverage structure. The model accurately differentiates between Equity

Value (Growth) and Credit Value (Solvency).

Comparative Analysis

Upon comparison with a linear model (e.g., Score = 0.5*Debt + 0.5*Coverage), CreditSentinel AI demonstrated higher model sensitivity. Under the linear model, high leverage is classified as moderately risky. However, CreditSentinel AI identified a non-linear risk escalation leading to default.

Limitations And Future Extensions

Data Heterogeneity and Dependence: The success of the current framework depends on the accuracy and reliability of the data ingestion provided by external sources, with particular emphasis on the Yahoo Finance API. While the Permissive Substitution Logic minimizes the effect of missing data, the current framework is still vulnerable to non-recurring accounting distortions or restatements that have not been captured in a timely manner in standardized data sources.

Integration of Unstructured Data: The current framework is largely quantitative in nature and fails to adequately capture qualitative and idiosyncratic risks such as litigation risks, regulatory risks, geopolitical risks, etc., which often act as key antecedents to financial distress. Future versions of the framework will utilize natural language processing tools to enable unstructured data ingestion and calculate a Sentiment Alpha score in addition to the current quantitative risk score.

Conclusion

CreditSentinel AI possesses significant potential to improve the workflow of institutional credit adjudication processes, while avoiding any replacement of human judgment. The framework helps to overcome the latency associated with human analyses, quantify governance risks that exist in emerging markets, and use non-linear modeling to deliver a submission-grade risk assessment in real-time. While it does not replace the role of the human credit adjudicator, it acts as an automated Risk Sentinel for the credit desk in a highly dynamic and evolving market environment.

References

1. EI Altman (1968) Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. The Journal of Finance 23: 589-609.
2. RC Merton (1974) On the pricing of corporate debt: The risk structure of interest rates. The Journal of Finance 29: 449-470.
3. T Chen, C Guestrin (2016) XGBoost: A scalable tree boosting system. in Proc. 22nd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining.
4. Reserve Bank of India (2019) Report of the Working Group on Introduction of Interest Rate Futures. RBI Bulletin.
5. J Rosenbaum, J Pearl (2013) Investment Banking: Valuation, Leveraged Buyouts, and Mergers and Acquisitions, 2nd ed. Hoboken NJ: Wiley.
6. Google (2014) Google Collaboratory.
7. R Aroussi (2014) yfinance: Yahoo! Finance market data downloader. Python package version 0.2.40.