Arm's Length Lending for the Thin-Filed Using Artificial Intelligence

Aditya Gurematti, Revansiddha Basavaraj Khanapure** & Tejasvi Varadaraj****

DOI: https://doi.org/10.5281/zenodo.21205154

ABSTRACT

This article reviews existing techniques and proposes new methods to evaluate credit risk in economies without formal credit systems but widespread mobile phone usage. Limited availability of formal financial data remains the primary drawback to credit access, yet AI-based analysis of digital behavior can help bridge this gap. Our review suggests that applying machine learning methods to classify behaviors such as gambling and alcohol-related spending as high risk may accurately capture credit risk, while expenditure on education, including school fee payments, may signal creditworthiness. Further, stable and long-term location patterns may serve as strong indicators of high credit quality. We propose using AI to discover hidden "non-linear" patterns, such as a combination of alcohol consumption and irregular phone charging, which may predict high credit risk but possibly escape human analysts. The proposed systems automate the search for these factors and create models that are robust to sparse data. Finally, we analyze the "Digital Utility Trap", where the fear of losing essential mobile access motivates borrowers to pay, offering a safe and scalable path for lending to the unbanked and thin-filed.

KEYWORDS: *Digital Collateral, Financial Inclusion, Thin-file Borrowers, Alternative Data, Machine Learning.*

1. Introduction

We present a unified framework for credit risk assessment that eliminates the necessity of centralized bureau data. In developing economies, the majority of adults are "thin-file" borrowers, possessing limited or no formal credit records. The World Bank Global Findex Database indicates rapid growth in individuals holding a mobile-money account only, as well as those holding both a bank (or similar financial institution) account and a mobile-money account. To quantify this growth: 40 percent of adults in Sub-Saharan Africa had a mobile money account as of 2024, up from 27 percent in 2021. In Latin America and the Caribbean, 37 percent of adults had a mobile money account as of 2024, up from 22 percent in 2021. Mobile-money accounts drive a substantial portion of this financial inclusion. Simultaneously, mobile phone ownership is high, reaching 86 percent of adults worldwide in 2024. Individuals generate continuous streams of digital activity every day.

Traditional banking infrastructure consistently fails to convert this ubiquitous digital footprint into usable credit signals. However, through advanced algorithmic intervention, this digital footprint can be leveraged to extend credit to thin-file borrowers who are systematically excluded by traditional banking systems (Berg et al., 2020). The absence of a traditional credit history represents an infrastructural failure, not a lack of underlying creditworthiness. This proposed system extracts risk signals from daily digital routines, using stable behavior as a proxy for financial reliability. We outline an architecture that combines alternative data, payment patterns, social networks, and machine learning to accurately assess credit risk.

The central contribution of this work is the introduction of behavioral stability metrics integrated with a built-in repayment incentive mechanism. We characterize the borrower's

* **Aditya Gurematti,** School of Natural Sciences and Mathematics, 800 West Campbell Road, The University of Texas at Dallas, Richardson, TX, USA 75080.

** **Revansiddha Basavaraj Khanapure,** Naveen Jindal School of Management, 800 West Campbell Road, University of Texas at Dallas, Richardson, TX, USA 75080.

*** **Tejasvi Varadaraj,** Eric Johnson School of Engineering, 800 West Campbell Road, The University of Texas at Dallas, Richardson, TX 75080.

reliance on mobile services as a natural deterrent to default. The following sections detail the data engineering pipelines required to process anonymized metadata and mobile-money logs. We explain how graph neural networks could capture homophily within social networks, and how supervised autoencoders might resolve missing alternative data. Finally, an implementation pathway for deploying this architecture is delineated, ensuring privacy and fairness.

2. The Economic Mechanism of Digital Collateral

Our framework is theoretically grounded in the economic theories of asymmetric information and commitment devices. Traditional credit markets rely on physical collateral to mitigate moral hazards. Physical assets signal wealth and provide recourse in default. Thin-file borrowers lack these assets, causing traditional models to reject them regardless of intent or capacity to repay. We advocate for substituting physical collateral with digital collateral, which we define as the quantifiable stability of an individual's digital life.

The mechanism driving dispersion in our envisioned behavioral variables is the heterogeneity in daily digital consumption. Variances in routine reflect underlying risk aversion and time preference, which remain the fundamental determinants of creditworthiness. A user maintaining a routinized geographic pattern and predictable communication schedule shows a low-variance lifestyle, which we argue translates directly into a lower probability of financial default. Conversely, highly erratic digital behaviors-including frequent geographic changes, irregular communication, and high-velocity small-denomination transactions-would signal a high-variance lifestyle associated with elevated credit risk.

High time-preference individuals heavily discount the future, prioritizing immediate consumption. This manifests as the immediate depletion of data balances or frequent, small-scale airtime borrowing. Low time-preference individuals ration data usage and maintain balances over longer horizons. By measuring these micro-behaviors, lenders could observe time preference, overcoming adverse selection without requiring extended credit histories.

We posit that the threat of digital exclusion could serve as a synthetic commitment device. This approach relies on the borrower's dependence on the mobile ecosystem. In emerging markets, mobile devices serve as the primary conduit for business, remittances, and communication. The utility derived from continuous access vastly exceeds the monetary value of a typical micro-loan. The threat of losing network access creates an overwhelming deterrent to default, potentially aligning the borrower's incentives with the lender's objectives perfectly.

3. Synthesis of Existing Empirical Evidence

Our envisioned architecture builds upon validated empirical successes. Djeundje et al. (2021) proves conclusively that alternative data sources effectively substitute for missing credit histories. They deliver substantial gains in predictive accuracy for thin-file borrowers, establishing the baseline viability of non-traditional variables in formal risk models. The specific mobile behavioral patterns yielding the highest predictive power would be isolated within our framework. Gambacorta et al. (2019) provide critical evidence from a Chinese financial technology firm, proving that machine-learning models fed with device usage logs and social footprints consistently outperform traditional demographic scorecards.

Tobback and Martens (2019) establish the predictive dominance of fine-grained payment details. Specific merchant categories, exact transaction timing, and recurrence patterns yield highly accurate risk profiles. We incorporate an adaptation of their methodology into our feature engineering process, where consistent payments to educational institutions would serve as a direct proxy for long-term investment behavior.

For relational data analysis, Muñoz-Cancino et al. (2022) prove that combining transaction networks and social connections significantly strengthens creditworthiness assessments. Their successful application of graph neural networks forms the basis of our hypothesized network analysis layer.

To address missing data, we turn to the methodology of Abdoli et al. (2021). They introduce the bagging of supervised autoencoders, successfully extracting robust latent features from highly imbalanced and incomplete datasets. Supervised autoencoders ensure the model remains stable despite noisy and sparse alternative data. Ayari et al. (2025) confirm hybrid machine learning models using alternative data currently dominate the credit-scoring literature. Davuluri et al. (2019) provides a successful real-world deployment analogue, showing that applying machine learning to behavioral data profitably expands credit access. This framework distinguishes itself by theoretically unifying these disparate, successful methodologies into a single, cohesive, and fully automated pipeline.

4. Proposed Feature Engineering and Data Pipeline

We detail a theoretical data engineering pipeline to ingest, clean, and transform raw telecommunication and mobile-money logs into predictive behavioral features. Lenders would need to establish secure data pipelines with operators and acquire explicit, opt-in consent from all users. Raw data sources would include mobile-money ledgers, call metadata, application usage categories, and cell-tower location traces. Three primary categories of engineered features are conceptualized: risk-signaling ratios, stability metrics, and network indicators.

The risky-spend ratio can be defined as total outbound mobile-money volume directed toward betting platforms and alcohol vendors, divided by total outbound volume. The literature establishes that high expenditures in these categories correlate strongly with default. Furthermore, the positive-signal ratio would track the volume of recurring payments directed toward essential services like school fees and water utilities, which directly indicate high financial discipline.

Furthermore, we introduce stability metrics to quantify digital collaterals. SIM tenure can be calculated as the integer number of months since initial network registration; long tenure indicates an established identity. Location stability would be assessed using the Shannon entropy of daily cell-tower connections over a rolling ninety-day window. A low entropy score would indicate predictable geographic routines, which strongly proxy for steady employment. Variance in the timing and frequency of outbound calls would also be measured, as stable daily routines exhibit low communication entropy.

5. Proposed Machine Learning Architecture

We outline a three-stage machine learning architecture to process engineered features: a graph representation layer, an autoencoder imputation layer, and a gradient-boosted classification layer. This combination addresses the unique properties of alternative data, including non-linear interactions and high rates of missing values.

First, the model would rely on constructing a directed graph where nodes represent users, and edges represent mobile-money transfers. The application of two layers of Graph Convolutional Networks (GCNs) is recommended to aggregate features from a user's direct connections. This would create a node embedding mathematically capturing the credit quality of the user's broader social network. It operationalizes homophily, capturing the reality that individuals connected to high-risk borrowers often show high-risk behaviors themselves.

Second, the autoencoder imputation layer draws upon Abdoli et al. (2021). Because alternative data streams suffer frequent interruptions, we envision employing an ensemble of supervised autoencoders to train the model to reconstruct the complete feature vector from corrupted inputs. The bottleneck layer would produce a low-dimensional latent representation of user behavior, remaining robust even when forty percent of input features are missing.

Third, the methodology involves concatenating the graph embeddings, autoencoder latent features, and engineered behavioral metrics into a single vector inputting into an XGBoost classifier. XGBoost naturally handles complex, non-linear interactions. It correctly classifies a user who shows occasional high-risk spending but maintains perfect location stability as low-risk-an interaction linear models strictly fail to capture. Feature discovery would be automated using SHapley Additive exPlanations (SHAP) values, implementing a recursive elimination loop to prune variables and eliminate human bias.

6. Policy Implications and Implementation Pathways

Deploying this AI Trust Score framework would require rigorous oversight to protect consumer privacy and ensure algorithmic fairness. The reliance on highly granular alternative data introduces ethical considerations that traditional bureau scoring avoids. We advise that financial institutions launch initial deployments within controlled regulatory sandboxes, allowing operators and lenders to share anonymized datasets under strict supervision.

A strict mandate for cryptographic anonymization is essential before data enters the pipeline. Lenders must never access raw communication content or specific geographic coordinates; the system should only ingest aggregated metadata. Continuous bias auditing is strongly recommended to prevent machine learning models from perpetuating existing societal inequalities. Institutions should use stratified cross-validation during training to test disparate impact across gender, income levels, and rural versus urban demographics.

Additionally, the use of SHAP values should be required for individual decision explainability. When the model rejects an applicant, the lender could provide a specific, behavior-based reason for the denial using the SHAP output, satisfying emerging regulatory requirements for algorithmic accountability. Governments and central banks could actively enhance this framework by artificially enriching positive data signals. Routing public subsidies or educational financial support through established mobile-money channels would directly improve the baseline credit scores of vulnerable populations. The combination of private algorithmic innovation and public policy integration maximizes the absolute impact of the digital collateral framework.

7. Conclusion

This paper presents a practical, scalable framework for credit-risk assessment in emerging markets characterized by high mobile penetration and absent credit bureaus. By synthesizing validated machine learning techniques, we illustrate how everyday digital habits could be transformed into concrete financial signals. Through the conceptual formalization of digital collateral, we hypothesize that behavioral stability can effectively replace physical assets in risk modeling. The borrower's fundamental dependence on mobile services could be harnessed as a built-in synthetic commitment device ensuring high repayment rates. With fair implementation, financial institutions can provide credit to millions of excluded adults. Ultimately, a lack of traditional credit history is no longer a barrier to accurate risk assessment and broad financial inclusion.

References:

1. M. Abdoli, M. Akbari, and J. Shahrabi, "Bagging supervised autoencoder classifier for credit scoring," arXiv preprint arXiv:2108.07800, 2021.
2. H. Ayari, R. Guetari, and N. Kraïem, "Machine learning powered financial credit scoring: A systematic literature review," Artificial Intelligence Review, vol. 59, p. 13, 2025.
3. V. B. Djeundje, J. Crook, R. Calabrese, and M. Hamida, "Enhancing credit scoring with alternative data," Expert Systems with Applications, vol. 163, 113766, 2021.
4. L. Gambacorta, Y. Huang, H. Qiu, and J. Wang, "How do machine learning and non-traditional data affect credit scoring? New evidence from a Chinese fintech firm," BIS Working Papers, no. 834, 2019.
5. R. Muñoz-Cancino, C. Bravo, S. A. Ríos, and M. Graña, "On the combination of graph data for assessing thin-file borrowers' creditworthiness," Expert Systems with Applications, vol. 188, 118809, 2022.
6. E. Tobback and D. Martens, "Retail credit scoring using fine-grained payment data," Journal of the Royal Statistical Society: Series A (Statistics in Society), vol. 182, no. 4, pp. 1227–1246, 2019.
7. S. Davuluri, R. García Franceschini, C. R. Knittel, C. Onda, and K. Roache, "Machine learning for solar accessibility: Implications for low-income solar expansion and profitability," NBER Working Paper No. 26178, 2019.
8. World Bank Group, The Global Findex Database 2021: Financial inclusion, digital payments, and resilience in the age of COVID-19, 2022.