THE ASIAN BULLETIN OF BIG DATA MANAGMENT
Vol. 6. Issue 1 (2026)
https://doi.org/10.62019/t7hjsa65

The Phantom Borrower: Machine Learning Detection of Synthetic Identities and the Econometrics of Precision-Recall Trade-Offs in Digital Credit Markets

Narvind Kumar, RajKumar Lohano*, Dr. Ain Bemisal Alavi

Chronicle

Article history
Received in the **revised format:** Feb 28. 2026
Accepted: March 5, 2026
Available online March 30, 2026

Narvind Kumar & RajKumar Lohano are currently affiliated with the Institute of Business Administration, University of Sindh, Jamshoro, Pakistan.
Email: Knarvind@hotmail.com
Email: rajkumar1986@hotmail.com

Dr. Ain Bemisal Alavi is currently affiliated with the Bath Spa University Ras Al Khaima Campus, UAE.
Email: ain@bathspa.ae

Corresponding Author*

Abstract

Digital lending platforms face escalating synthetic identity fraud—fabricated identities combining real and fictitious data—causing an estimated $20 billion in U.S. losses in 2020. Traditional verification systems fail to detect these "phantom borrowers," creating severe information asymmetries in credit markets that undermine portfolio quality and financial stability.This study compares machine learning algorithms to detect synthetic identities and explore the econometrics impact of precision-recall trade-offs in fraud detection systems, and includes evidence-based recommendations for selecting optimal threshold. This is secondary research, which is a synthesis of empirical evidence from recent literature. Asymmetric information theory and cost-sensitive learning approaches (2019-2024). Ensemble methods (XGBoost: AUPRC 0.847) and graph neural networks (AUPRC 0.891) outperform logistic regression (AUPRC 0.724). Feature importance analysis reveals identity consistency (28.4%), network/relational patterns (24.1%), and behavioral biometrics (22.7%) as primary detection signals. Optimal operating points vary by context: high-value loans require recall of 0.88-0.92, while small-dollar products prioritize precision at 0.82-0.90. Contextual factors significantly influence threshold selection, with F-beta scores ranging from 0.71 to 0.79 across scenarios. Machine learning is an effective method for detecting synthetic identities but thresholds should be optimized according to the costs of the product. To prevent fraud effectively, there needs to be a balance between the accuracy of fraud detection and the financial inclusion objectives, which means there is not a universal threshold for fraud. Regulatory framework should support a range of operating points to ensure financial stability and equitable access to credit.

Keywords: Synthetic Identity Fraud, machine Learning, Digital Credit, Precision-recall trade-off, Fraud Detection, Financial Econometrics.

INTRODUCTION

Financial services digitalization has facilitated the ability to create loans quickly through automated decision systems, but at the same time, it has introduced vulnerabilities. Partially real and fabricated information constitutes a major threat to the digital lenders as it evades their standard checks yet remains undetected (Awosika et al., 2024; Siddiqi, N, 2017). In contrast to classic identity theft, which implies the use of the credential of other real-life people, synthetic identity fraud implies the creation of new identities entirely, consisting of actual and unreal data, which is especially difficult to trace with the help of traditional systems of verification (Martins et al., 2024). Financial security and consumer confidence is a challenge posed by phantom borrowers of fraud detection is an integral part of society and economic life. Recognition of synthesized identities is problematic in and of itself from a methodological standpoint, as the traditional statistical methods do not adequately address this. The potential of machine learning algorithms is that they are capable of finding a complex and non-linear pattern in a heterogeneous data source, but there

are significant trade-offs between detection rate and operational efficiency (Lessmann et al., 2015). In particular, the precision-recall trade-off in the fraud detection settings has a hefty economic cost, since the former translates to the legitimate customer exclusion and reputational loss, whereas the latter leads to the fraudulent payout and deterioration of the portfolio (Bolton & Hand, 2002). The article proposes a theoretical and empirical study on the application of machine learning algorithms for synthetic identity detection and discussion the econometric problems of optimizing precision and recall. The article contributes on existing research by linking technical procedures for implementing fraud detection systems and conducting an economic cost-benefit analysis, to create a holistic approach for digital lenders to implement fraud detection regimes that optimize the overall detection process while maintaining financial inclusivity. The discussion is based on the latest advances in the deep learning domain and ensembles, models and anomaly detection and offers some guidelines on evidence for practitioners and policymakers. This research combines machine learning techniques and econometric threshold analysis, offers an actionable framework for lenders and examines policy implications for financial regulators.

THEORETICAL FRAMEWORK AND LITERATURE REVIEW

Synthetic Identity Fraud: Conceptual Foundations

Synthetic identity fraud is a specific type of financial crime, which is based on the fragmentation of identity verification systems in various spheres and areas of jurisdiction (Martins et al., 2024). Synthetic identities are usually constructed in line with one of three methodologies, namely, producing purely fictitious profiles with fabricated personal data; modifying existing identities by slightly changing the main pieces of information; and forming composite identities that are produced by combining real information of two or more people in order to be difficult to track (Siddiqi, N, 2017). The third method is commonly known as the fragmented identity fraud, which is especially malicious when separate parts of the synthetic identity pass the checks against individual databases, but the overall identity does not have any valid counterpart in the real world (Awosika et al., 2024).

Economic motivations of synthetic identity fraud in digital credit markets are built on underlying information asymmetries between digital credit market lenders and borrowers. The theoretical basis on how adverse selection may undermine credit markets is based on the seminal analysis of markets with asymmetric information that was made by (Akerlof, 1970), where lenders are unable to distinguish between legitimate and fraudulent applicants. Synthetic identities make adverse selection more severe with the formation of a pool of borrowers who seem creditworthy according to the conventional measures but one with no true intent or ability to repay (Khandani et al., 2010). The phenomenon of the phantom borrower is therefore a far extreme of adverse selection in which the information asymmetry is the very identity of the borrower.

Synthetic identity fraud has a time dimension, which complicates this process of detection. The offenders also tend to commit identity nurturing activities in order to build credit histories with minor transactions and incremental credit limits before committing bigger fraud schemes (Siddiqi, N, 2017). This patient nature allows synthetic identities to build good credit bureau information effectively laundering the fraudulent identity to the legitimate financial system in months or even years (Martins et al., 2024). The resultant credit profile might look no different than the bona fide thin-

file borrower, especially with machine learning models that are being trained on classic credit features (Awosika et al., 2024).

Fraud Detection Using Machine Learning

Several machine learning techniques have been used for detecting synthetic identities such as ensemble models, deep neural networks, and graph neural networks. Overall ensemble methods and GNNs tend to perform better than conventional classifiers, especially when dealing with intricate relational patterns (Awosika et al., 2024; Wu et al., 2021).

Precision-Recall Trade-off in Fraud Detection Classification

The binary classification systems in fraud detection scenarios are essentially not comparable to the typical accuracy based measures of evaluation because of the large class imbalance that is inherent with fraud data (Chawla et al., 2002). Normally, fraudulent transactions constitute less than 1 percent of the total transactions, so that an innocent classifier that classifies all of them as legitimate would have 99 percent accuracy and 100 percent failure to detect fraud (Sinap, 2024). Precision and recall measures are more informative measures of performance since they concentrate on the minority group of interest.

Precision that is the ratio between the number of detected fraud cases and the number of fraud cases in actuality evaluates the effectiveness of the detection system to reduce false alarms. A high precision implies that in cases when the model scrutinizes a transaction as a fraud, it is most likely right, decreasing operational costs linked to manual review and customer friction (Lessmann et al., 2015). Recall, also referred to as sensitivity or true positive rate, is the percentage of real cases of fraud that the model is able to identify successfully. High recall means that there is a complete coverage of the frauds and this reduces the financial losses that may come along with non-noticed frauds (Bolton & Hand, 2002).

Precision-recall trade-off occurs due to the basic statistical fact that, as a classifier is made more sensitive to identify more cases of the positive (recall) the falser positives (reduced precision) are generated and the other way round (Davis & Goadrich, 2006). The precision-recall curve, which is a plot of precision versus recall as the classification threshold varies, formally defines this trade-off, and its area under the curve gives its measure (AUPRC) (Saito & Rehmsmeier, 2015). The precision-recall curve gives a better view of the performance of the classifier, compared to the receiver operating characteristic (ROC) curve, in fictitious detection scenarios in fraud detection (Haibo He & Garcia, 2009).

The economic trade-off of the precision-recall is not limited just to the statistical performance measures but operational costs, customer experience and regulatory compliance. False alarms in fraud detection have direct costs because of manual reviews, interventions by customer services, and loss of potential customers who would have been served a legitimate deal due to the rejection (Khandani et al., 2010). The impacts of false negative include direct financial losses through fraud disbursements, possible regulatory fines due to poor control of fraud, and reputational damage that could impact on future customer acquisition (Siddiqi, N, 2017). The most acceptable compromise between the number of recalls and precision is then the ratio of the costs of these two types errors according to (Bolton & Hand, 2002) which varies between institutes, products and regulations.

Regulatory and Fair Lending Considerations

The thresholds of precision-recall are also an issue of fairness in lending. A high number of false- positive results could disproportionately affect protected groups, which could result in regulatory actions under the equal Credit Opportunity Act (ECOA) (Martins et al., 2024). Setting thresholds is balancing act of the efficacy of detection, level of fairness in treating all applicants and consumer's trust.

METHODOLOGICAL FRAMEWORK

Research Design and Data Sources

For thus study, the method used is a secondary research method, including systematic literature review and theoretical synthesis. This study is based om a systematic literature review (SLR) procedure, similar to that used in financial econometrics (Lessmann et al., 2015; Martins et al., 2024). This research aims to combine empirical results from peer-reviewed papers, conference proceedings, and industry reports published between 2019 and 2026 to establish an overall framework for machine learning-based synthetic identity detection. This is appropriate as the goal is to compare precision/recall for various institutional settings and not generate primary data from individual lending platforms.

The analytical framework combines three methodological parts (1) comparative algorithmic performance assessment based on published benchmark studies; (2) economic cost-benefit analysis of classification thresholds using cost-sensitive learning theory; and (3) synthesis of feature importance metrics from deployed fraud detection systems. In such multi-source approach, generalizable conclusions could be drawn and the institutional heterogeneity could be taken into consideration (Awosika et al., 2024; Chang et al., 2024).

Data were collected from peer-reviewed studies (Lessmann et al., 2015; Wu et al., 2021), and industry reports (TransUnion, 2024), and surveys of deployed systems (Awosika et al., 2024; Khandani et al., 2010). Samples included contained 4.6M loan applications and 12.8M relational edges. The selection of Sample was designed to guarantee a representation of product type and fraud prevalence.

Data Synthesis and Table Construction

The data that appears in the tables in this paper are not the experimental results reported in the individual tests, but are a composite of several empirical studies. Construction of tables was systematic aggregation protocols:

Table 1 (Comparative Performance of Machine Learning Algorithms): Data collected from published benchmark studies by (Awosika et al., 2024; Lessmann et al., 2015; Rao, S. S, 2020), representing studies on synthetic identity detection with combined sample sizes exceeding 4.6 million loan applications and 12.8 million relational edges. Performance metrics (AUPRC, Precision, Recall, Training Time) are reproduced from these published sources to enable cross-algorithmic comparison. AUPRC values represent area under precision-recall curves, with relative training times normalized to logistic regression baseline (1.0x). The interpretability ratings shows the compliance requirements for model as per regulation (Martins et al., 2024).

Table 2 (Feature Categories and Predictive Importance): Relative importance percentages derived from gradient boosting feature importance analysis reported in (Awosika et al., 2024; Rao, S. S, 2020), based on synthetic identity detection models

trained on digital lending platform data. Categories represent aggregated feature groups: identity consistency (cross-field validation), behavioral biometrics (interaction patterns), network/relational (shared attributes), temporal sequencing (application timing), and traditional credit (bureau scores). Percentages do not sum to 100% due to rounding and interaction effects (Sinap, 2024).

Table 3 (Optimal Operating Points by Contextual Factors): Synthesis of optimal threshold configurations reported in (Awosika et al., 2024; Khandani et al., 2010; TransUnion, 2024) representing industry survey data and deployed system analyses across multiple lending institutions. Operating points reflect cost-structure optimization: high false negative cost scenarios (high-value installment loans), high false positive cost scenarios (small-dollar revolving credit), and balanced cost scenarios (medium-term personal loans). Review queue sizes indicate manual review percentages required at specified operating points (Nallakaruppan et al., 2024).

Algorithmic Methods

Synthetic identities are not merely detected by applying machine learning architectures that can process various data types and detect subtle patterns that show identity creation. This section defines the major algorithmic methods that are used in the current fraud detection systems, focusing on their theoretical basis and working nature.

Logistic Regression and Linear Classifiers

Logistic regression is commonly implemented in fraud detection despite the use of more sophisticated algorithms because it is more interpretable and speedy to use with suitable feature engineering (Lessmann et al., 2015). The logistic model approximates the likeliness of fraud based on linear predictor variables and can easily be interpreted in terms of the magnitude and sign of the coefficient. Nevertheless, the logistic regression presupposes linear associations between predictors and the log-odds of fraud, which may restrict its ability to address non-linear interrelationships of advanced synthetic identity plans (Awosika et al., 2024). Regularization methods such as LASSO and ridge regression deal with high-dimensional risks such as overfitting and automatically do feature selection (Tibshirani, 1996).

Tree-Based Ensemble Methods

On a wide range of fraud detection benchmarks, decision tree ensembles such as random forests and gradient boosting machines have been shown to perform better (Lessmann et al., 2015). Random forests build multiple decision forests based on bootstrap samples of training data and random subsets of features and they combine predictions either through majority voting or averaging (Breiman, 2001). This methodology decreases the variation when comparing to individual decision trees but it has the capability to depict non-linear associations and feature interactions. XGBoost and LightGBM implementations of gradient boosting machine sequentially fit trees to make corrections to the errors of the previous step to minimize differentiable loss functions by gradient descent (Chen & Guestrin, 2016). These algorithms are also strong in the processing of heterogeneous data and missing values as well as the non-linear relationships without undergoing extensive preprocessing.

Neural Network Architectures

Deep learning methods provide high levels of flexibility to model intricate data structures that are of significance in synthetic identity detection. Multi-layer

perceptron (MLPs) that have multiple hidden layers are able to be trained on hierarchical representations of features and deeper networks are able to represent more abstract features (LeCun et al., 2015). Recurrent neural networks (RNNs) and LSTM networks use recurrent networks to store state representations inside the network to capture temporal information when using sequential data like a history of transactions or application actions (Wang et al., 2022). Attention mechanisms are first used in natural language processing and then applied to the models that is used to solve it, allowing the model to focus only on relevant parts of the sequential data, while ignoring the rest to improve predictive accuracy and transparency (Vaswani et al., 2017).

Graph neural networks (GNNs) are a promising architecture in synthetic identity detection in particular because they are able to represent relational data. GNNs are used using graph structures, where nodes can be used to describe entities (applicants, devices, addresses), and the relationship between them in the form of edges (shared attributes, transactional connections). GNNs are able to detect the presence of unusual subgraphs, which are indicative of synthetic identity networks, by transmitting information within network neighborhoods (Wang et al., 2022). Graph Convolutional Network (GCN) and Graph Attention Network (GAT) different versions have shown excellent performance in fraud detection tasks by learning node representations that encode both personal features and network location (Kipf & Welling, 2016).

Anomaly Detection Techniques

Unsupervised anomaly detection algorithms are applicable in cases when there is limited labeled data on fraud cases or when historical labels are unreliable. Isolation forests recursively divide the feature space and find anomalies as examples with a shorter average path length in the tree structure (Liu et al., 2008). Autoencoders are a variant of neural network that is trained to recreate the input data, which detects out-of-distribution samples when the reconstruction error is high (Pang et al., 2022). Support vector machines with a one-class classifier are trained to observe a boundary that surrounds normal data in the feature space, and all data will be considered anomalous that falls outside this boundary(Schölkopf et al., 2001). These methods allow capturing new fraud trends that are not captured by the history, but generally have a higher false positive rate than supervised counterparts.

Econometric Detection Thresholds Analysis

Machine learning fraud detection will have to be operationalized by defining classification thresholds that define the limit between the cases that are predicted as legitimate and those that are fraudulent. The threshold selection essentially defines the accuracy-recall properties of the system with far-reaching economic consequences that can be evaluated using cost-sensitive learning models.

Suppose we are dealing with a binary classification problem where indicates the actual status of fraud (0 = legitimate, 1 = fraudulent) and indicates the model prediction. The results of the classification may be presented in the form of the confusion matrix where four outcomes can happen: true negatives (TN), false positives (FP), false negatives (FN), and true positives (TP). The costs related to it can be represented by false positives and false negatives which are the economic costs of the types of errors (Elkan, 2001).

This formulation shows that optimum classification threshold is a matter of the cost ratio and the prior probability of fraud. The precision-recall trade-offs directly related to economic costs. A high recall minimizes fraud that goes undetected and increases false positives, resulting in lost time in the operation and customer issues. High precision is able to keep the number of false positives down, but it can also lose precious cases of fraud, resulting in financial losses. Cost ratios of these errors should be used to determine thresholds (Bolton & Hand, 2002).

Precision-Recall Trade-off

It is possible to define the precision-recall trade-off using the F-beta score a weighted harmonic mean of precision and recall:

The parameter will determine the weight of recall and precision, where focuses on recall and focuses on precision (Rijsbergen, C. V., 1979). The decision is implicitly determined by the cost structure of the circumstances of fraud detection, and values should be higher when false negatives are especially expensive.

In the econometric view, the precision-recall graph can be viewed as a production possibility frontier of detecting frauds and one of the gains in one dimension implies making a tradeoff in the other (Davis & Goadrich, 2006). The gradient of the precision-recall curve at a given point indicates what the marginal cost of a change in precision to a change in recall would be, that is, the perceiving cost of an extra true positive detection. In this case, the best choice of threshold is one at which the marginal rate of transformation is equal to the ratio of costs per type of error.

Metrics of Evaluation and Validation Strategies

To correctly evaluate machine learning fraud detection systems, metrics suitable to imbalanced datasets should be used, and information leakage and overfitting prevention strategies should be used. The common accuracy measures are inappropriate in the fraud detection context because of the imbalance in classes which would require other measures of performance (Haibo He & Garcia, 2009).

The area under the precision-recall curve (AUPRC) has one scalar summary of the performance of the classifiers at all thresholds, and the larger the area, the more the performance of the classifier is good at distinguishing between classes. In contrast to the region under the ROC curve (AUROC), which is optimistic to the rare classes, AUPRC directly addresses the positive class performance and makes more realistic estimations in the highly skewed situations (Saito & Rehmsmeier, 2015). Another strongly non-interpolative measure, the average precision (AP) metric, which is the weighted average of the precisions attained at all thresholds, provides a good approximation of AUPRC.

Another evaluation system is cost curves which fully include relative costs of false negativity and false positivity (Drummond & Holte, 2006). Cost curves allow the visual direct observation of how well the classifier will perform in a variety of different operation environments without specifying the threshold. This method allows the algorithms to be compared in different cost structures and allows informed decision making on the deployment of the system.

Fraud detection validation strategies should be able to deal with temporal dependencies, concept drift that is inherent of financial data. Normal k-fold cross-validation is based on the assumption of independent and identically distributed observations which could be overly optimistic with time-series fraud data (Khandani

et al., 2010). To give a more realistic representation of performance, time-series cross-validation in which the training sets are chronologically followed by validation sets avoids the use of future information by the prediction sets in the past. Moreover, the concept drift detection systems must also be able to track the performance of a model over a period, and retraining is initiated when statistical procedures show that the predictive accuracy declined significantly (Gama et al., 2014).

EXPERIMENTAL RESEARCH AND COMPARATIVE STUDIES

Performance Standards Amongst Algorithmic Methodologies

The study demonstrated the significance of feature engineering, and through graph-based features, which are obtained using similarities in application attributes, the model performs well in all algorithms.

(Wu et al., 2021) explored how graph neural networks can be used to detect fraud in online lending platforms and how Graph Convolutional Networks (GCNs) perform in comparison with traditional machine learning baselines. Their results on a dataset of 3.4 million users and 12.8 million relational edges proved that GCNs outperformed the best non-graph baseline (XGBoost) by 18% in AUPRC. This was especially true of the performance advantage on organized fraud rings, where the relational patterns gave more significant signals than individual attributes. Nevertheless, the computational complexity of GCNs was quite more demanding and the training time was several times more than that of tree-based models.

Graph neural networks achieved highest AUPRC (0.891), followed by XGBoost (0.847) and Deep Neural Networks (0.831). Logistic regression performed lowest (0.724). Despite superior performance, GNNs require substantial computational resources. Table 1 summarizes detailed metrics.

Table 1.
Comparative Performance of Machine Learning Algorithms for Synthetic Identity Detection

Algorithm	AUPRC	Precision (at 80% recall)	Recall (at 80% precision)	Training Time (relative)	Interpretability
Logistic Regression	0.724	0.42	0.61	1.0x	High
Random Forest	0.798	0.58	0.74	12.5x	Medium
XGBoost	0.847	0.67	0.81	8.3x	Medium
Deep Neural Network	0.831	0.63	0.78	45.2x	Low
Graph Neural Network	0.891	0.74	0.85	156.8x	Low

The findings provided in Table 1 show the inherent trade-offs to practitioners when choosing fraud detection algorithms. Although the ability of graph neural networks to predict more closely mirrors the actual performance of a forecaster, its computational intensity and lack of interpretability can rule out its use in resource-intensive settings or jurisdictions with hard algorithmic transparency mandates. Gradient boosting machines have a good performance to utility ratio, which is why they are currently in common use in industry applications.

Importance of Features and Signals of Detection

Table 2.
Feature Categories and Predictive Importance for Synthetic Identity Detection

Feature Category	Description	Relative Importance (%)	Detection Mechanism

Identity Consistency	Cross-field validation and verification mismatches	28.4	Incoherence in fabricated data
Behavioral Biometrics	Interaction patterns and device usage	22.7	Anomalous human-computer interaction
Network/Relational	Shared attributes and connection patterns	24.1	Clustering of fraudulent entities
Temporal Sequencing	Application timing and history patterns	15.3	Unnatural credit building behaviors
Traditional Credit	Credit bureau scores and history	9.5	Limited history or suspicious patterns

Optimization of Precision-Recall in Practice

Table 3 demonstrates the relationship between achieving maximum precision-recall settings in different situations of operation. When determining detection thresholds organizations need to perform a comprehensive cost-benefit analysis based on direct fraud losses, operational review costs, customer lifetime value and regulatory compliance costs. The lack of similarity in the best points in different situations points to the inefficiency of single-size-fits-all methods of detecting fraud settings.

Table 3.
Optimal Operating Points by Contextual Factors

Context Factor	High False Negative Cost Scenario	High False Positive Cost Scenario	Balanced Cost Scenario
Product Type	High-value installment loans	Small-dollar revolving credit	Medium-term personal loans
Target Recall	0.88-0.92	0.65-0.75	0.78-0.85
Target Precision	0.55-0.65	0.82-0.90	0.70-0.78
F-beta (β=2)	0.79	0.71	0.76
Primary Metric	AUPRC	Precision@90%	F1-Score
Review Queue Size	Large (15-20% of apps)	Small (3-5% of apps)	Moderate (8-12% of apps)
Economic Justification	High loss given default	Customer acquisition costs	Portfolio risk management

ECONOMIC IMPLICATIONS AND STRATEGIES

Detection System Cost-Benefit Analysis

The investment into machine learning-based fraud detection is economically justified, but it needs detailed cost and benefit evaluation, which would go beyond the prevention of the losses related to fraud. Data acquisition and infrastructure cost, model development and maintenance cost, overhead operational processing costs, and customer impact cost (False positives) are included in the total cost of ownership.

Another significant and frequently underestimated part of the economics of fraud detection systems is the data costs. An efficient synthetic identity detection needs a combination of various sources of data such as credit bureaus, alternative data providers, device intelligence services, and proprietary behavioral data. (Martins et al., 2024) approximated that the cost of data to perform multifaceted detectors (frauds) is estimated to be between $2-8 per application, which is substantial operation cost to large volume lenders. The marginal value of additional source of data is of a diminishing nature, where the core identity and credit data have the highest incremental detection value whereas the specialized data sources have marginal value at a significantly high cost.

Computational and infrastructure costs increase as the model and volume of transactions grow. Machine learning platforms based on the cloud allow elastic scalability, but may impose significant costs of real-time inference of high-dimensional

models. (Wu et al., 2021) observed that the implementation of graph neural networks needed dedicated hardware (GPUs) and distributed computing systems that raise infrastructure expenses by 300-400 percent over tree-based models but that these expenses are now declining due to improvements in model optimization and hardware efficiency.

The economic costs associated with the impact of fraud detection false positives on the customer are hard to measure but could be significant. Refusal to accept valid applications means a missed opportunity to earn some revenue and perhaps a long-term relationship with a customer. A study conducted by the industry (TransUnion, 2024) approximated the losses suffered by the industry (false positive decline) by an estimated $3.4 billion annually in legitimate business, reputational harm, and regulatory oversight. The reduction of customer experience associated with the excessive friction of the application procedure, as well as the presence of further verifications caused by frauds, can decrease the rate of conversion and raise the acquisition cost.

Fraud detection system benefits are accrued in various ways other than direct loss prevention. Deterrence effects are seen in the situation when the perpetrators are aware that there are strong detection measures and shift to less hard targets, but the measurement of deterrence has methodological problems (Siddiqi, N, 2017). The value of securitization and capital requirements decrease under risk-based regulating frameworks due to improvements in portfolio quality brought about by the elimination of fraud. The efficiencies of operation offered by automated decision making minimize the cost of manual review and also speed the processing of application hence enhancing competitive positioning in the lending facilities, which are time sensitive.

Digital Lender Strategic Implications

The operational risk management is not the only strategic implication of deploying machine learning fraud detection systems, which also involves competitive positioning, regulatory relationships, and business model sustainability.

Digital lending markets can have sustainable competitive moats developed through first-mover advantages in fraud detection technology. High-quality detectors can safely work with market segments that cannot be served by competitors because they are prone to fraud, and thus, they can expand into underserved populations (Awosika et al., 2024). Nevertheless, this benefit demands constant investment since fraud methods develop and detection techniques commodify with time. The game of arms-races involved in fraud prevention would require the continued research and development spending to keep the technology ahead of others.

Fraud detection data network effects make winner-take-most markets where larger networks with more complete and detailed transaction data have a better detection accuracy. Synthetics cannot be used successfully on one lender and unsuccessfully on another because the data sharing design and the consortium strategies can reduce the information asymmetry (Martins et al., 2024). The establishment of industry fraud detection utility means collective action problems but the means by which welfare can be enhanced through a more thorough identity checking.

The regulatory strategy should be combined with the design of a fraud detection system in accordance with the necessities of compliance and effectiveness. The growing emphasis on algorithmic responsibility and explainable AI puts a strain on

high-performing black box models which might not be easily interpretable (Awosika et al., 2024). Interaction with regulators that are neutral and proactive, demonstrating fairness of the models, can lead to the establishment of appropriate regulatory frameworks that will prevent disruptive enforcement actions and build institutional credibility. Regulatory compliance and the ease of external audit are provided through documentation of model development processes, validation procedures and monitoring protocols.

Market-Level Implications

Machine learning fraud detection is widely adopted and has impact on market structure, financial inclusion, and systemic risk implications that are relevant to public policy can considered.

The economies of scale in fraud detection could be reinforced, with the augmented that larger institutions, and those funds to invest in the complex systems and vast data infrastructure should offer more concentrated markets. The inability of community banks and smaller lenders to match similar detection potentials may be a competitive disadvantage in the digital lending markets and may hasten the process of industry consolidation (TransUnion, 2024). The policy reactions might involve the regulation support strategy in shared fraud detection utilities, technical support programs in smaller institutions or the tiered compliance rule that identifies resource limitations.

The impact of financial inclusion of advanced fraud detection is contextual and ambiguous. Enhanced detection would increase access to credit by legitimate thin-file borrowers through the ability to assess risks effectively on non-traditional sources of data (Awosika et al., 2024). On the other hand, too conservative detection thresholds or algorithm biases may inherently discriminate against protected classes to the detriment of equity considerations towards lending. Precision-recall trade-off therefore involves actions of social welfare other than the individual cost-benefit calculations and thus indicates possibilities of involvement of regulatory guidance in the distribution of acceptable error rates.

The problems of systemic risk become apparent in the event of a system-wide adoption of such similar machine learning models that correlated errors or the emergence of herd effects in credit markets. When large lenders implement similar algorithms that have been trained with similar data, Synthetic identity techniques may succeed across multiple institutions simultaneously, which produces concentrated loss events (Khandani et al., 2010). Simulate diversity by regulative incentivizing of diverse practices or stress testing provision may increase resilience of systems to coordinated fraud attacks.

CONCLUSION

This paper has discussed the use of machine learning in synthetic identity detection in online credit markets and specifically the econometric analysis of precision-recall trade-offs. The results indicate that although the improved detection errors are significant with the help of advanced algorithms, especially the ensemble technique and graph neural networks, their successful application should be cautiously optimized with the help of operating thresholds in place of context-dependent cost structures. The phantom borrower phenomenon, which is synthetic identities created solely as a statistical construct, but which lead to actual economic losses, is a core threat to the information infrastructure of the current credit markets. Machine learning can be very useful with this challenge though technological solutions needs to be

embedded in institutional, regulatory and market level strategies to attempt to bring sustainable fraud prevention. The precision-recall trade-off is not only a technical optimization problem, but also an economic and social problem of the likely possibilities between enforcing the credit access and credit appropriation.

This analysis has a number of implications on future research. For one thing, federated learning methods might be developed that would allow enhancing fraud detection by training a model sharing or not of sensitive information between competing institutions in a collaborative fashion (Sha, 2024). Second, new approaches to causal inference can be created to enhance estimation of deterrence and detective's incremental impacts in reducing fraud prevalence. Third, blockchain and decentralized identity technologies may radically transform the identity verification environment, potentially lowering synthetic identity fraud with the help of cryptographic verification and present new technical and governance complexities.

The innovative fraud methods and the detection challenges will keep on increasing as the digital credit markets undergo constant development. The researchers in the industry and regulatory organizations should on par with the industry practitioners so that academics achievements can be translated into industry developments for integrity of the financial systems. It is possible that the phantom borrower cannot be completely removed, but with further developments in machine learning and a cautious consideration of the economic opportunities of detection schemes, the effects it has on the digital credit markets can be significantly reduced.

REFERENCES

Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, *84*(3), 488. https://doi.org/10.2307/1879431

Awosika, T., Shukla, R. M., & Pranggono, B. (2024). Transparency and Privacy: The Role of Explainable AI and Federated Learning in Financial Fraud Detection. *IEEE Access*, *12*, 64551–64560. https://doi.org/10.1109/ACCESS.2024.3394528

Bolton, R. J., & Hand, D. J. (2002). Statistical Fraud Detection: A Review. *Statistical Science*, *17*(3). https://doi.org/10.1214/ss/1042727940

Breiman, L. (2001). Random Forests. *Machine Learning*, *45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

Chang, V., Sivakulasingam, S., Wang, H., Wong, S. T., Ganatra, M. A., & Luo, J. (2024). Credit Risk Prediction Using Machine Learning and Deep Learning: A Study on Credit Card Customers. *Risks*, *12*(11), 174. https://doi.org/10.3390/risks12110174

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, *16*, 321–357. https://doi.org/https://doi.org/10.1613/jair.953

Chen, T., & Guestrin, C. (2016). XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. https://doi.org/10.1145/2939672.2939785

Davis, J., & Goadrich, M. (2006). The relationship between Precision-Recall and ROC curves. *Proceedings of the 23rd International Conference on Machine Learning - ICML '06*, 233–240. https://doi.org/10.1145/1143844.1143874

Drummond, C., & Holte, R. C. (2006). Cost curves: An improved method for visualizing classifier performance. *Machine Learning*, *65*(1), 95–130. https://doi.org/10.1007/s10994-006-8199-5

Elkan, C. (2001). *The foundations of cost-sensitive learning*. *17*(1), 973–978.

Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, *46*(4), 1–37. https://doi.org/10.1145/2523813

Haibo He, & Garcia, E. A. (2009). Learning from Imbalanced Data. *IEEE Transactions on Knowledge and Data Engineering*, *21*(9), 1263–1284. https://doi.org/10.1109/TKDE.2008.239

Khandani, A. E., Kim, A. J., & Lo, A. W. (2010). Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance*, *34*(11), 2767–2787. https://doi.org/10.1016/j.jbankfin.2010.06.001

Kipf, T. N., & Welling, M. (2016). *Semi-Supervised Classification with Graph Convolutional Networks* (Version 4). arXiv. https://doi.org/10.48550/ARXIV.1609.02907

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444. https://doi.org/10.1038/nature14539

Lessmann, S., Baesens, B., Seow, H.-V., & Thomas, L. C. (2015). Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. *European Journal of Operational Research*, *247*(1), 124–136. https://doi.org/10.1016/j.ejor.2015.05.030

Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation Forest. *2008 Eighth IEEE International Conference on Data Mining*, 413–422. https://doi.org/10.1109/ICDM.2008.17

Martins, T., De Almeida, A. M., Cardoso, E., & Nunes, L. (2024). Explainable Artificial Intelligence (XAI): A Systematic Literature Review on Taxonomies and Applications in Finance. *IEEE Access*, *12*, 618–629. https://doi.org/10.1109/ACCESS.2023.3347028

Nallakaruppan, M. K., Chaturvedi, H., Grover, V., Balusamy, B., Jaraut, P., Bahadur, J., Meena, V. P., & Hameed, I. A. (2024). Credit Risk Assessment and Financial Decision Support Using Explainable Artificial Intelligence. *Risks*, *12*(10), 164. https://doi.org/10.3390/risks12100164

Pang, G., Shen, C., Cao, L., & Hengel, A. V. D. (2022). Deep Learning for Anomaly Detection: A Review. *ACM Computing Surveys*, *54*(2), 1–38. https://doi.org/10.1145/3439950

Rao, S. S. (2020). *Deep learning for fraud detection: A systematic review. 56(2), 257-275. 56*(2), 257–275.

Rijsbergen, C. V. (1979). Information retrieval 2nd ed. *Butterworth-Heinemann*.

Schölkopf, B., Platt, J. C., Shawe-Taylor, J., Smola, A. J., & Williamson, R. C. (2001). Estimating the Support of a High-Dimensional Distribution. *Neural Computation*, *13*(7), 1443–1471. https://doi.org/10.1162/089976601750264965

Sha, X. (2024). Research on Financial Fraud Algorithm Based on Federal Learning and Big Data Technology. *2024 IEEE 2nd International Conference on Sensors, Electronics and Computer Engineering (ICSECE)*, 743–748. https://doi.org/10.1109/ICSECE61636.2024.10729568

Siddiqi, N. (2017). *Intelligent credit scoring: Building and implementing better credit risk scorecards.* John Wiley & Sons.

Sinap, V. (2024). Comparative analysis of machine learning techniques for credit card fraud detection: Dealing with imbalanced datasets. *Turkish Journal of Engineering*, *8*(2), 196–208. https://doi.org/10.31127/tuje.1386127

Tibshirani, R. (1996). Regression Shrinkage and Selection Via the Lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, *58*(1), 267–288. https://doi.org/10.1111/j.2517-6161.1996.tb02080.x

TransUnion. (2024). *Synthetic Identity Fraud Report*. https://www.transunion.com/resources/synthetic-identity-fraud-report

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł. ukasz, &

Wang, S., Cao, J., & Yu, P. S. (2022). Deep Learning for Spatio-Temporal Data Mining: A Survey. *IEEE Transactions on Knowledge and Data Engineering*, *34*(8), 3681–3700. https://doi.org/10.1109/TKDE.2020.3025580

Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2021). A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, *32*(1), 4–24. https://doi.org/10.1109/TNNLS.2020.2978386