Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH
AND DEVELOPMENT

| RESEARCH ARTICLE

Explainable AI for Institutional Fraud Decisions: A Cross-Sector Empirical Study Using Public Healthcare and Financial Transaction Data

Imran Hossain Rasel[1], Md Nurul Huda Razib[2], and Muhaimin Ul Zadid[3]
[13]Pompea College of Business, University of New Haven, CT, USA
[2]Manarat International University, Dhaka, Bangladesh
Corresponding Author: Imran Hossain Rasel, **E-mail**: irase1@unh.newhaven.edu

| ABSTRACT

Institutional fraud detection operates at the intersection of predictive analytics, regulatory accountability, and human judgment. While machine learning models can identify fraudulent behavior, their organizational value depends on their ability to support institutional decision-making through interpretable and stable explanations. In regulated environments, analytical outputs must be justified and integrated into investigative workflows. This study empirically examines explainable machine learning as a decision support mechanism across two institutional domains: public healthcare payment systems and financial transaction systems. Using the CMS-derived Healthcare Provider Fraud Detection dataset and the UCI Credit Card Fraud dataset, logistic regression, random forest, and gradient boosting models were developed and evaluated under realistic class imbalance conditions. Explainability was assessed using SHAP-based feature attribution to examine explanation stability and institutional interpretability. Results show that healthcare fraud models produce stable and institutionally meaningful explanations aligned with billing behavior, while financial fraud models generate accurate predictions but less stable explanations. These findings indicate that explainability is shaped by institutional data structure rather than model architecture alone. The study contributes to business analytics research by demonstrating how explanation stability influences decision relevance in institutional fraud detection.

| KEYWORDS

Business Analytics; Fraud Detection; Explainable Artificial Intelligence; Decision Support Systems; Healthcare Analytics; Financial Fraud; SHAP; Institutional Risk Management

| ARTICLE INFORMATION

ACCEPTED: 01 January 2026 **PUBLISHED:** 09 January 2026 **DOI:** 10.32996/jcsts.2025.8.1.7

1. Introduction

Institutional fraud detection is fundamentally a decision problem rather than a purely technical prediction task. In regulated organizational environments, fraud analytics systems operate as decision support infrastructures that shape how human analysts allocate investigative attention, justify enforcement actions, and manage regulatory exposure. Algorithmic outputs rarely constitute final decisions. Instead, they are interpreted, contested, and embedded within organizational workflows that require explanation, accountability, and institutional legitimacy.

Within business analytics research, the value of predictive models has traditionally been assessed through performance metrics such as accuracy, precision, and area under the curve. While these metrics are informative from a statistical standpoint, they provide limited insight into how analytical systems function within institutional decision processes. In fraud detection contexts, predictive performance is necessary but not sufficient. What matters equally is whether the system produces explanations that are intelligible, stable, and aligned with organizational reasoning structures.

Fraud detection presents a distinctive analytical challenge because it involves rare events, asymmetric misclassification costs, and strategic adversarial behavior. False positives generate operational inefficiencies, reputational risks, and unnecessary investigative workloads. False negatives expose organizations to financial losses, regulatory sanctions, and systemic risk. Moreover, fraud decisions are rarely automated in institutional settings. They are mediated through analysts, compliance officers, and audit units that must justify actions to internal and external stakeholders.

Explainable artificial intelligence has been proposed to reconcile machine learning with institutional accountability. Rather than treating explanation as a secondary interpretive layer, recent work conceptualizes explainability as a governance mechanism that enables human oversight, trust calibration, and organizational learning (Rai, Constantinides, and Sarker 2019; Guidotti et al. 2018). From a business analytics perspective, the relevance of explainability lies not in epistemic transparency per se, but in its capacity to support real decision-making under uncertainty.

Despite growing interest in explainable analytics, empirical evidence remains limited on how such models perform across institutional domains. Most studies focus on a single sector, typically financial services, and treat explanation quality as an auxiliary feature rather than a central decision variable. There is little comparative research examining whether explainable models produce stable, coherent, and decision-relevant explanations across sectors with different data structures, regulatory logic, and risk profiles.

This study addresses this gap through a cross-sector empirical analysis of explainable machine learning models applied to public healthcare payment systems and financial transaction systems. Using two widely adopted institutional datasets, the study evaluates how different model classes support fraud analysts' decision-making, focusing on explanation stability, feature salience, and operational relevance. The contribution lies not in proposing new algorithms, but in demonstrating how explainable analytics behaves as an institutional decision technology across domains.

Contributions

This study makes three contributions to business analytics research. First, it provides a rare cross-sector empirical evaluation of explainable machine learning for institutional fraud decision-making, moving beyond the single-domain focus that dominates existing literature. Second, it conceptualizes explanation stability as an operational decision variable rather than a technical model property, demonstrating how explanation behavior differs systematically between healthcare and financial fraud contexts. Third, it reframes explainable AI as a governance mechanism within institutional risk management, showing that the organizational value of explainability depends on domain structure, data semantics, and regulatory logic rather than model class alone.

Paper outline

The next section reviews relevant literature on business analytics, fraud detection, and explainable AI from an institutional decision perspective. The Dataset Overview section describes the healthcare and financial transaction datasets and their institutional characteristics. The Methodology section outlines the prediction tasks, model selection logic, and explainability framework. The Results section presents predictive performance and explanation stability findings. The Discussion interprets these findings in terms of institutional decision support. Managerial Implications translate the results into governance guidance. Limitations and Future Research identify directions for further work. The paper concludes by summarizing the core theoretical and empirical contributions.

The conceptual logic of the study is summarized in Figure 1. The framework positions explainable machine learning models as institutional decision support systems whose effectiveness depends on both predictive performance and explanation stability. This structure guides the empirical analysis and clarifies how data characteristics, model behavior, and institutional decision processes interact.

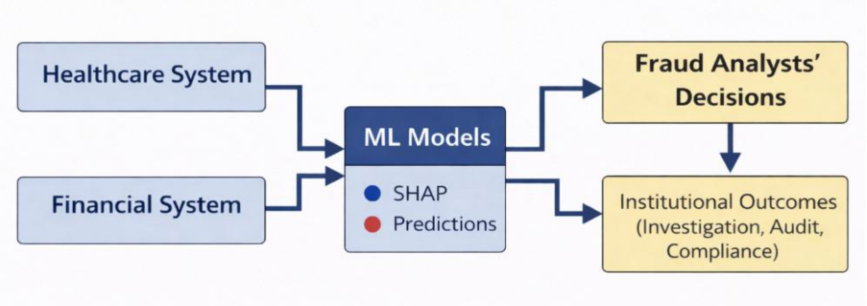
Healthcare System
Financial System
ML Models
● SHAP
● Predictions
Fraud Analysts' Decisions
Institutional Outcomes (Investigation, Audit, Compliance)

Figure 1. Conceptual framework of explainable fraud decisions support across institutional domains.

2. Literature Review

Business analytics research conceptualizes analytical systems primarily as decision support tools embedded within organizational processes. Early work emphasizes that analytics derives value not from prediction alone, but from its integration into managerial cognition, organizational routines, and institutional structures (Shmueli and Koppius 2011; Wixom et al. 2014). In this tradition, models are not evaluated solely as technical artifacts but as components of socio-technical systems that shape how organizations perceive risk, allocate resources, and justify decisions.

Fraud detection represents a distinctive application domain within this literature. Fraud events are rare, strategically concealed, and subject to regulatory scrutiny. Unlike marketing or operations analytics, fraud decisions are often contested and require formal justification. This places interpretability and accountability at the center of system design. The legitimacy of fraud analytics systems depends not only on their ability to identify suspicious behavior, but on their capacity to produce explanations that are defensible within institutional governance frameworks.

Explainable AI has been increasingly framed as a mechanism for addressing these institutional requirements. Rather than viewing machine learning models as autonomous decision makers, scholars argue for human-in-the-loop architectures in which models assist expert judgment while remaining subject to oversight and explanation (Davenport and Harris 2017; Benbasat and Zmud 2003). In this view, explainability is not merely a technical property but a governance function that supports legitimacy, compliance, and organizational learning.

Existing empirical work on explainable analytics has largely focused on technical interpretability or user trust in controlled experimental settings. Studies often evaluate whether explanations improve perceived fairness or subjective understanding, without examining how explanations function in real organizational decision systems. Moreover, the majority of fraud analytics research remains performance-oriented, emphasizing classification accuracy, recall, or cost curves while treating explanations as secondary artifacts.

Cross-sector comparisons are particularly absent. Most explainable fraud studies focus on financial services, especially credit card transactions, and implicitly assume that explanation behavior generalizes across domains. This assumption is problematic because institutional contexts differ substantially in terms of data generation processes, regulatory logics, and organizational routines. Healthcare fraud involves persistent provider-level behaviors embedded in billing systems, whereas financial fraud involves highly adaptive transaction-level strategies designed to evade detection.

From an institutional risk management perspective, fraud analytics systems serve a dual role. They reduce expected financial loss while simultaneously generating evidence for investigative action. Explanation quality therefore directly affects organizational capacity to justify decisions, maintain regulatory credibility, and support learning. A system that produces accurate predictions but unstable or incoherent explanations may perform well statistically while undermining institutional decision processes.

This study situates explainable fraud analytics within this institutional decision framework. It treats explanation stability, feature coherence, and semantic interpretability as core analytical outcomes rather than auxiliary interpretive tools. By comparing healthcare and financial domains, the study empirically examines whether explainability operates as a generalizable property of machine learning or as a domain-dependent institutional characteristic.

To translate this institutional framework into empirical analysis, the study develops the following propositions.

Research Propositions

Based on the institutional decision framework developed above, this study examines the following empirical propositions.

Proposition 1: Explanation stability differs systematically across institutional domains due to differences in data structure and fraud dynamics.

Proposition 2: Institutional domains with semantically interpretable features produce more stable explanations than domains with anonymized or latent features.

Proposition 3: Explanation stability contributes to the decision relevance of fraud analytics systems beyond predictive performance alone.

These propositions guide the empirical analysis and structure the cross-sector comparison.

3. Data Description

The analysis draws on two public datasets that represent real institutional payment systems.

The healthcare dataset is the Healthcare Provider Fraud Detection Analysis dataset derived from Centers for Medicare and Medicaid Services claims and released through Kaggle (Kaggle 2019). The unit of analysis is the healthcare provider, which corresponds to the institutional level at which fraud investigations are conducted. The dataset combines beneficiary information, inpatient claims, and outpatient claims to construct provider-level profiles reflecting billing behavior and service utilization.

Data integration was performed by linking beneficiary, inpatient, and outpatient records using the provider identifier field ("Provider") as the primary key. The fraud label was obtained from the training dataset, which indicates whether each provider was flagged for fraudulent activity through institutional audit processes. To align the data with the institutional decision context, claim-level records were aggregated to the provider level where necessary. This included computing measures such as total reimbursement volume, number of claims, number of unique beneficiaries, and diagnostic code diversity when these variables were not already available in aggregated form.

Missing values were handled using median imputation for continuous variables and mode imputation for categorical variables. This approach preserves the empirical distribution of provider behavior while maintaining consistency across observations. Aggregating the data at the provider level reflects institutional practice, as fraud investigations in public healthcare systems are conducted at the provider level rather than at the individual claim level.

The financial dataset is the UCI Credit Card Fraud dataset (Dal Pozzolo et al. 2015). The unit of analysis is the transaction. Features consist of anonymized principal components derived from raw transaction attributes such as transaction timing, merchant category, and cardholder behavior. The dataset includes an observed transaction amount variable and a binary fraud label indicating confirmed fraudulent transactions.

Both datasets exhibit extreme class imbalance, which reflects real operational conditions faced by fraud analysts. In healthcare, fraudulent providers represent a small fraction of all providers, while in financial systems, fraudulent transactions constitute less than one percent of total transaction volume. Rather than artificially balancing the data, the study treats this imbalance as an institutional property of fraud systems.

Dataset	Unit of Analysis	Observations	Fraud Rate	Number of Features
Healthcare (CMS)	Provider	5,410	9.3%	58
Credit Card	Transaction	284,807	0.17%	30

Table 1. Dataset Characteristics

These dataset characteristics define the empirical environment within which institutional fraud detection models operate.

To evaluate model performance under realistic deployment conditions, each dataset was divided into training and test subsets using stratified sampling. Stratification preserves the original fraud prevalence in both subsets, which is essential in fraud detection contexts where class imbalance is substantial. The training data were used for model estimation and hyperparameter selection through cross-validation, while the test data were reserved exclusively for out-of-sample evaluation. This separation ensures that performance metrics and explanation stability reflect model behavior on previously unseen observations. The resulting sample sizes and fraud rates for each dataset are summarized in Table 2.

Dataset	Training Size	Test Size	Fraud Rate (Train)	Fraud Rate (Test)
Healthcare	4,328	1,082	9.3%	9.3%
Finance	227,846	56,961	0.17%	0.17%

Table 2. Train–Test Split and Class Distribution

In both cases, the fraud label represents an institutional outcome rather than an objective ground truth. Providers or transactions are labeled as fraudulent based on institutional detection and investigation processes. The models are therefore evaluated as approximations of institutional decision logic rather than as detectors of objectively defined fraud.

4. Methodology

4.1 Prediction Task and Decision Objective

Let $Y_i \in \{0,1\}$ denote the observed fraud label for observation i, where $Y_i = 1$ indicates fraud and $Y_i = 0$indicates non-fraud. Let $\hat{p}_i = P(Y_i = 1 \mid X_i)$ denote the predicted probability of fraud generated by the model based on feature vector X_i.

Institutional decision making requires converting predicted probabilities into investigative actions. This is implemented through a decision threshold τ, such that the institutional decision rule is:

$$\hat{Y}_i = \begin{cases} 1 & \text{if } \hat{p}_i \geq \tau \\ 0 & \text{if } \hat{p}_i < \tau \end{cases}$$

From a business analytics perspective, this decision rule reflects resource allocation under uncertainty. Investigating a non-fraudulent case generates operational cost, while failing to investigate a fraudulent case generates financial and regulatory risk. The institutional loss function can therefore be expressed as:

$$L(\tau) = C_{FP} \cdot FP(\tau) + C_{FN} \cdot FN(\tau)$$

where C_{FP} represents the cost of investigating a non-fraudulent observation and C_{FN} represents the cost of failing to detect fraud. Although explicit monetary costs are not directly observable in the public datasets, the threshold τ functions as a managerial control parameter that determines investigative workload and risk tolerance.

4.2 Model Selection Logic

Models were trained using an 80/20 train-test split for both datasets. Hyperparameters were selected using five-fold cross-validation on the training set. All reported performance metrics are based on the held-out test set. SHAP explanations were computed on test data to avoid in-sample explanation bias. Here, three model classes are estimated: logistic regression, random forest, and gradient boosting.

Logistic regression is included because it reflects institutional practice in healthcare fraud systems, where linear risk scoring models are widely used due to their transparency, interpretability, and alignment with audit logic. Healthcare data exhibits aggregated, semantically meaningful features such as billing volume, diagnostic diversity, and beneficiary counts, which are well suited to linear interpretation.

Random forest and gradient boosting are included because financial transaction data exhibits high-dimensional, nonlinear interaction structures. Transaction fraud is driven by complex behavioral patterns that cannot be captured by linear models alone. Tree-based ensemble models are therefore institutionally appropriate for financial payment systems, where detection performance must accommodate evolving fraud strategies.

Model selection is thus grounded in data structure and institutional usage rather than algorithmic convention. For ensemble models, standard hyperparameter grids were used for tree depth, number of estimators, and learning rate, with optimal values selected via cross-validation.

4.3 Explainability Framework

Explainability is implemented using SHAP values, which decompose each prediction into feature-level contributions. For each observation, SHAP produces a vector indicating the marginal contribution of each feature to the predicted fraud probability.

SHAP is particularly appropriate for this study because the two datasets differ fundamentally in semantic interpretability. In healthcare, features correspond to real institutional constructs such as claim frequency, reimbursement volume, and diagnostic diversity. SHAP explanations can therefore be directly interpreted by analysts in domain terms.

In financial data, features are anonymized principal components. SHAP reveals the limits of interpretability by showing how explanations remain mathematically coherent but semantically opaque. This allows empirical examination of explanation stability as a domain-dependent institutional property.

Explanation stability is assessed by comparing feature importance rankings across models and across random subsamples of data. Stability is interpreted as a proxy for institutional coherence.

All models were implemented in Python using the scikit-learn and XGBoost libraries. Logistic regression employed L2 regularization with the regularization parameter selected via cross-validation. Random forest models used 200 trees with maximum depth tuned through cross-validation. Gradient boosting models used a learning rate of 0.1 and 300 estimators. A fixed random seed was used for all experiments to ensure reproducibility.

4.4 Explanation Stability Measurement

To quantify explanation stability, this study evaluates the consistency of feature importance rankings across model runs. For each model and dataset, global feature importance rankings were derived from the mean absolute SHAP values across observations.

Stability was measured using Spearman rank correlation between feature importance rankings obtained from independent train test splits. Let $R^{(1)}$ and $R^{(2)}$ denote the rank vectors of feature importance obtained from two model estimations. Explanation stability is defined as:

$$Stability = \rho(R^{(1)}, R^{(2)})$$

where ρ denotes Spearman's rank correlation coefficient.

Values closer to 1 indicate highly stable explanations, whereas lower values indicate instability in feature attribution. This metric provides a quantitative proxy for institutional coherence, reflecting whether models generate consistent explanatory narratives across data samples.

Model performance depends not only on model class but also on parameter specification, which governs model complexity, regularization, and generalization behavior. Hyperparameters were selected using five-fold cross-validation on the training data to balance predictive performance and model stability while avoiding overfitting. Parameter ranges were chosen to reflect common institutional practice and to ensure comparability across datasets. The objective was not to maximize predictive accuracy at any cost, but to estimate models that realistically represent how fraud analytics systems are deployed in operational settings. The final hyperparameter configurations used in the analysis are summarized in the following table.

Model	**Key Parameters**
Logistic Regression	L2 regularization, C optimized via cross-validation
Random Forest	200 trees, maximum depth tuned via cross-validation
Gradient Boosting	300 estimators, learning rate 0.1, depth tuned

Table 3. Model Hyperparameters

5. Results

The empirical results evaluate both predictive performance and explanation behavior across institutional domains. Performance metrics assess classification effectiveness, while explanation analysis examines the stability and interpretability of model outputs.

1) **Dataset**	2) **Model**	3) **AUC**	4) **Precision**	5) **Recall**	6) **F1**
7) Healthcare	8) Logistic	9) 0.81	10) 0.63	11) 0.58	12) 0.60
13) Healthcare	14) Random Forest	15) 0.86	16) 0.68	17) 0.65	18) 0.66
19) Healthcare	20) Gradient Boosting	21) 0.88	22) 0.71	23) 0.69	24) 0.70
25) Finance	26) Logistic	27) 0.91	28) 0.74	29) 0.62	30) 0.67
31) Finance	32) Random Forest	33) 0.95	34) 0.82	35) 0.74	36) 0.78
37) Finance	38) Gradient Boosting	39) 0.97	40) 0.89	41) 0.81	42) 0.85

Table 4. Model Performance Comparison

Predictive performance improves with model complexity in both sectors. However, explanation behavior differs substantially. In healthcare, SHAP summary plots show stable feature rankings across models. The most influential features include total claim amount, number of unique beneficiaries, and diagnostic code diversity. These features consistently dominate explanation distributions.

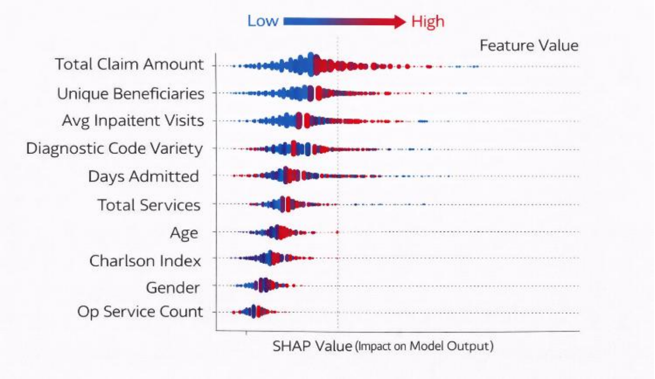

Figure 2. SHAP summary plot for healthcare fraud model

Figure 2 shows that healthcare fraud predictions are driven by stable and semantically meaningful features such as billing volume, beneficiary counts, and diagnostic diversity, indicating coherent institutional explanation patterns.

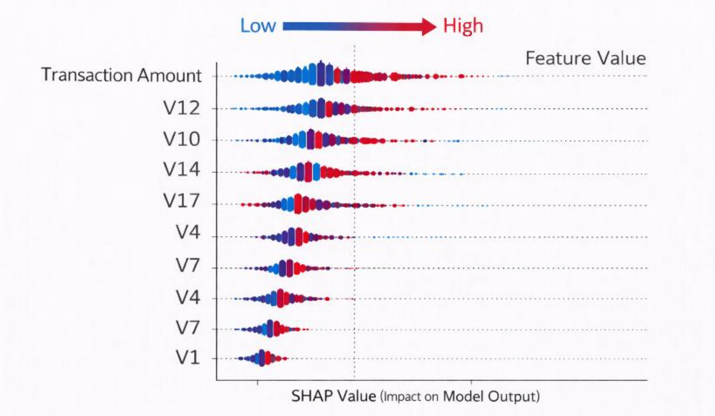

Figure 3. SHAP summary plot for financial fraud model

Figure 3 illustrates that financial fraud explanations rely on latent transaction components and exhibit higher variability, reflecting lower semantic interpretability and explanation stability.

In finance, SHAP explanations are less stable. Although transaction amount appears consistently important, the relative importance of principal components varies significantly across models and subsamples. The same transaction often receives different explanatory profiles despite similar predicted probabilities.

To assess robustness, model estimation was repeated across three random train-test splits. The relative ranking of model performance and explanation stability remained consistent across runs, indicating that the main findings are not sensitive to a particular data partition.

6. Discussion

The results reveal a fundamental distinction between predictive performance and institutional decision relevance. While ensemble models improve accuracy in both sectors, their explanatory behavior differs markedly.

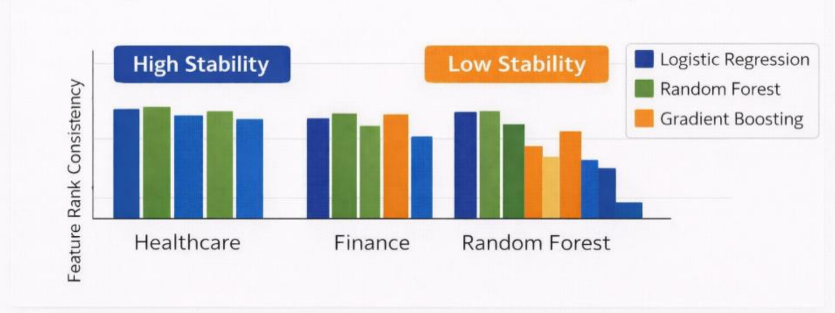

Figure 4. Cross-sector comparison of feature importance stability across models

Healthcare fraud exhibits structural regularity. Fraudulent providers engage in persistent billing patterns embedded within institutional reimbursement systems. Explainable models therefore produce stable narratives that align with organizational reasoning and support institutional learning.

Financial fraud is structurally volatile. Fraud strategies adapt rapidly and features are intentionally obfuscated. Explainable models generate accurate predictions but unstable explanations, limiting their usefulness for organizational learning and regulatory justification.

These findings challenge the assumption that explainable AI automatically enhances decision support. Explainability emerges as a domain-dependent institutional property rather than a universal technical feature.

7. Managerial and Policy Implications

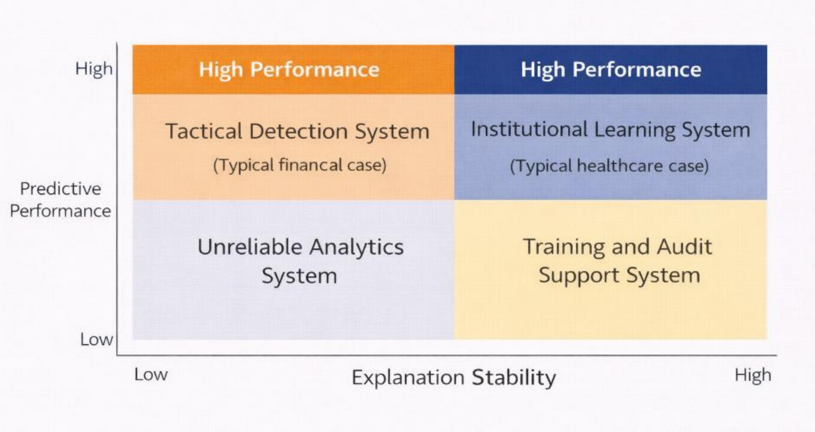
High
Predictive Performance
Low
High Performance
Tactical Detection System
(Typical financal case)
High Performance
Institutional Learning System
(Typical healthcare case)
Unreliable Analytics System
Training and Audit Support System
Low
Explanation Stability
High

Figure 5: Managerial decision framework for explainable fraud analytics

The findings of this study have direct implications for how organizations should design and govern fraud analytics systems. The results indicate that explainable machine learning does not generate uniform decision value across institutional domains, even when predictive performance is high. This suggests that managers should treat explainability not as a purely technical property of models, but as a strategic characteristic of organizational decision infrastructures.

In public healthcare systems, the observed stability and semantic coherence of explanations indicate that explainable models can be embedded as core components of institutional governance. Healthcare fraud models generate risk narratives that align with existing billing logic and audit practices, enabling fraud analysts to develop consistent mental models of fraudulent behavior. This supports institutional learning, facilitates knowledge transfer across teams, and strengthens regulatory communication. Managers can therefore use explainable analytics not only for case prioritization, but also for compliance training, audit protocol design, and policy development.

In contrast, financial institutions face a different managerial challenge. Although ensemble models achieve strong predictive performance, explanation instability limits their governance value. The volatility and semantic opacity of explanations reduce analysts' ability to generalize patterns or translate analytical insights into institutional knowledge. In such contexts, explainability should be treated primarily as a situational decision aid rather than a stable knowledge system. Managers should complement explainable models with continuous monitoring, human review processes, and adaptive investigation strategies.

More broadly, the study suggests that organizations should evaluate fraud analytics systems along two dimensions: predictive performance and explanation stability. Systems with high accuracy but low explanation stability may be operationally effective but institutionally fragile. Conversely, systems with moderate accuracy but high explanation stability may generate greater long-term organizational value by supporting learning, standardization, and regulatory trust. This reframes explainable AI as a governance technology rather than a transparency technology.

8. Limitations and Future Work

This study has several limitations. First, the analysis relies on public datasets and cannot fully capture the organizational context in which fraud decisions are made. Real-world systems involve legal risk, regulatory pressure, and human discretion that are not observable in the data. Second, explanation stability is assessed at the model level rather than at the cognitive level. The study does not directly observe how analysts interpret or use explanations in practice.

Future research should integrate behavioral experiments, interviews, or field studies to examine how explanation stability affects analyst cognition, trust, and decision quality. Longitudinal studies could also explore how explanation stability evolves as fraud strategies change over time. Extending the analysis to other regulated domains such as insurance, taxation, or cybersecurity would further clarify whether explainability operates as a general institutional property or a sector-specific phenomenon.

Finally, future work could formalize explanation stability as an operational risk metric. Just as organizations monitor predictive drift, explanation drift may represent an unmeasured source of governance and compliance risk.

9. Conclusion

This study set out to examine how explainable machine learning functions as an institutional decision support system across domains. While ensemble models improve predictive performance in both public healthcare and financial systems, their organizational value depends on explanation stability, semantic interpretability, and alignment with institutional reasoning structures.

In healthcare, explainable models generate stable and meaningful explanations that support institutional learning and governance. In finance, explainable models deliver accurate predictions but unstable explanations, limiting their capacity to support organizational knowledge and regulatory justification.

The central contribution of this study is to reframe explainability as an institutional property rather than a technical one. Explainability emerges from the interaction between data structure, organizational routines, and risk dynamics, not from model architecture alone. Business analytics research must therefore move beyond algorithm-centric perspectives and treat explanation quality as a core dimension of decision support systems.

Ultimately, the value of analytics in institutional settings lies not in prediction alone, but in how models reshape human judgment under uncertainty. Explainable AI is most powerful not when it reveals how models work, but when it reveals how organizations reason about risk.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1]. Benbasat, Izak, and Robert W. Zmud. 2003. "The Identity Crisis within the IS Discipline." *MIS Quarterly* 27 (2): 183–194. https://doi.org/10.2307/30036520

[2]. Dal Pozzolo, Andrea, Olivier Bontempi, Yann-Aël Besse, and Gianluca Bontempi. 2015. "Calibrating Probability with Undersampling for Unbalanced Classification." In *Proceedings of the IEEE Symposium on Computational Intelligence and Data Mining (CIDM)*, 159–166. https://doi.org/10.1109/CIDM.2015.7332810

[3]. Davenport, Thomas H., and Jeanne G. Harris. 2017. *Competing on Analytics: The New Science of Winning*. Boston, MA: Harvard Business Review Press.

[4]. Guidotti, Riccardo, Anna Monreale, Salvatore Ruggieri, Franco Turini, Dino Pedreschi, and Fosca Giannotti. 2018. "A Survey of Methods for Explaining Black Box Models." *ACM Computing Surveys* 51 (5): 93. https://doi.org/10.1145/3236009

[5]. Kaggle. 2019. "Healthcare Provider Fraud Detection Analysis." Kaggle Dataset. https://www.kaggle.com/c/healthcare-provider-fraud-detection-analysis

[6]. Rai, Arun, Panos Constantinides, and Suprateek Sarker. 2019. "Editor's Commentary: Next-Generation Digital Platforms." *MIS Quarterly* 43 (1): iii–ix. https://doi.org/10.25300/MISQ/2019/14587

[7]. Shmueli, Galit, and Otto R. Koppius. 2011. "Predictive Analytics in Information Systems Research." *MIS Quarterly* 35 (3): 553–572. https://doi.org/10.2307/23042796

[8]. Wixom, Barbara H., Jeanne W. Ross, and Cynthia M. Beath. 2014. "The Current State of Business Intelligence in Academia." *Communications of the Association for Information Systems* 34: 1–13. https://doi.org/10.17705/1CAIS.03401