



AI-Based Analytics for Financial Risk Assessment and Anomaly Detection in Complex Financial Data Systems

Sadia Zaman¹; Binayan Dey²;

[1]. Management Information Systems, College of Business, Lamar University, USA;
Email: zaman.sadia1311@gmail.com

[2]. Assistant Manager, Systems & IT, Chittagong Stock Exchange Ltd, Bangladesh.
Email: binayan.dey@gmail.com

Doi: [10.63125/224h4a32](https://doi.org/10.63125/224h4a32)

This work was peer-reviewed under the editorial responsibility of the IJBEI, 2023

Abstract

A modern financial institution watches an unrelenting stream of transactions, positions, prices, and counterparty signals, and somewhere within that stream sit the rare events, a fraud, a limit breach, a building loss, a contagion, that matter most and reveal themselves least. This study is about the machinery organizations build to catch those events: AI-based analytics for financial risk assessment and anomaly detection in complex financial data systems. Rather than asking whether a particular algorithm detects anomalies, it asks a broader, more organizational question, what makes the whole detection-and-response apparatus effective? To answer it, the study models financial risk-management effectiveness as an outcome of six capabilities: risk data aggregation and lineage, AI risk-modeling capability, anomaly detection performance, early-warning and signal timeliness, model risk governance and validation, and an integrating construct of detection–response integration. Evidence came from a structured five-point Likert survey completed by 158 valid respondents out of 176 distributed, an 89.8% valid response rate, drawn from risk analysts and managers, quantitative and model developers, data scientists, compliance and audit staff, and risk-technology specialists, of whom 69.0% worked directly with risk models or anomaly detection. The data were analyzed with descriptive statistics, Cronbach's alpha, Pearson correlation, multiple regression, a system maturity index, and a risk-analytics capability priority matrix. Every capability was rated in the high band, led by financial risk-management effectiveness at a mean of 4.26 and anomaly detection performance at 4.19; reliability was strong, with alpha from 0.83 to 0.94. All correlations were positive and significant, the strongest being $r = 0.81$ between detection–response integration and effectiveness. The regression model explained 74.1% of the variance in effectiveness, $R^2 = 0.741$, adjusted $R^2 = 0.729$, $F(6,151) = 71.94$, $p < 0.001$, with detection–response integration the leading predictor, $\beta = 0.33$, ahead of anomaly detection performance, $\beta = 0.26$, and AI risk-modeling capability, $\beta = 0.22$. The central lesson is that detection and response are one system: catching an anomaly matters only if the organization can assess and act on it, and the institutions that manage risk well are those that have wired sensing, modeling, and response into a single vigilant loop rather than a chain of disconnected tools.

Keywords

Financial risk assessment; Anomaly detection; AI analytics; Model risk governance; Early-warning systems

INTRODUCTION

Risk in a financial institution is rarely loud. It hides in the ordinary flow of business, in a transaction that looks like thousands of others until it is not, in a position that drifts past a limit before anyone notices, in a correlation that quietly strengthens until it becomes contagion. The defining challenge of financial risk management is therefore one of attention: how to watch an enormous, fast-moving, and complex data stream closely enough to catch the rare events that carry outsized consequences, without being paralyzed by the far larger volume of events that do not (Chandola et al., 2009). Artificial intelligence has become central to meeting that challenge, because the scale and subtlety of modern financial data exceed what human analysts or fixed rules can monitor (Cao, 2022). Surveys of deep learning in finance and of the field's research landscape document how rapidly these methods have spread across risk, trading, and detection applications (Ozbayoglu et al., 2020), (Goodell et al., 2021), and dedicated treatments now set out machine learning for finance as a discipline in its own right (Dixon et al., 2020). This study examines the systems institutions build for that purpose, AI-based analytics for financial risk assessment and anomaly detection, and asks what makes them effective as a whole.

The core computational task can be framed simply. An anomaly detector assigns each observation a risk score and raises a flag when that score crosses a threshold. This framing conceals every hard decision in risk analytics: what to measure, how to learn normal behavior, where to set the threshold that trades missed events against false alarms, and how to act once a flag is raised. Anomaly detection also sits alongside, and feeds into, the quantitative measurement of risk itself, whose two workhorse measures are value-at-risk and expected shortfall. Anomaly detection and risk measurement are complementary: detection surfaces the unusual events that risk models may not yet reflect, while risk measures quantify the exposure those events imply (Jorion, 2007). This study treats both as facets of a single analytics capability rather than as separate technical exercises.

Figure 1: AI-Based Analytics for Financial Risk Assessment and Anomaly Detection



The study's central argument is that the effectiveness of this capability is decided less by the sophistication of any single detector than by how well the surrounding system is assembled (Amir & Chapal, 2022; Abdur & Iftekhhar, 2021). A powerful anomaly model built on poorly aggregated data will miss what it cannot see; an accurate detector whose alerts no one can interpret or act on will not reduce loss; a well-governed model that signals too late will warn of events that have already happened (Leo et al., 2019). Research on AI in finance has repeatedly found that value comes from an integrated capability, data, models, governance, and response working together, rather than from technology in isolation (Mikalef et al., 2018). The same holds with particular force in risk, where the cost of a fragmented system is measured in losses that were detectable but not detected, or detected but not addressed (Hull, 2018). Two considerations make this an urgent subject. First, financial data systems have grown genuinely complex (Hasan & Uddin, 2022; Mohiul & Badrul, 2022): high volume, high velocity, heterogeneous sources, and dense interconnection mean that risks propagate through networks in ways that defeat siloed monitoring (Battiston et al., 2016). Second, the AI methods now

used to watch these systems are themselves a source of risk, opaque, drift-prone, and difficult to validate, which is why model risk governance has become inseparable from detection performance (Federal Reserve, 2011). This study brings these threads together, framing anomaly detection, risk measurement, model governance, and response as components of one capability whose contribution to risk-management effectiveness can be estimated.

Background of the Study

The turn to AI-based risk analytics has been driven by a mismatch between the complexity of financial data and the capacity of traditional tools to monitor it. Rule-based systems, long the mainstay of transaction monitoring and limit control, are transparent and cheap but brittle: they catch only what they were told to look for, generate high false-positive rates, and are easily evaded by behavior that shifts to stay just inside the rules (West & Bhattacharya, 2016). As data volumes grew and threats adapted, institutions turned to machine-learning methods that learn normal behavior from data and flag deviations from it, a family of techniques with deep roots in the anomaly-detection literature (Ahmed et al., 2016).

These methods have matured quickly. Unsupervised detectors such as isolation forests and local-outlier-factor methods identify anomalies without labeled examples, which suits a setting where confirmed events are rare (Liu et al., 2008). Supervised and deep-learning models, where labels exist, have improved detection of complex patterns in fraud and credit risk (LeCun et al., 2015). Streaming methods now score events in near real time and adapt as distributions shift, addressing the drift that plagues static models (Gama et al., 2014). Yet the same maturation has surfaced new difficulties: advanced detectors are often opaque, their alerts hard to explain to investigators and regulators, and their behavior hard to validate, problems that the governance and explainability literatures have only begun to resolve (Rudin, 2019). This background frames the study's motivation. The detection methods are, individually, well developed and widely available. What remains uncertain is how institutions assemble them, together with data infrastructure, governance, and response processes, into a coherent system, and how much such a system actually improves risk-management effectiveness. The study addresses that uncertainty by treating data aggregation, modeling capability, detection performance, warning timeliness, governance, and detection-response integration as measurable capabilities whose contributions can be estimated.

Problem Statement

The problem this study addresses is that investment in AI-based risk analytics does not automatically translate into better risk management. Institutions acquire sophisticated detection models yet continue to suffer losses that were, in principle, detectable, because data was fragmented and incomplete, because alerts arrived too late or in unmanageable volume, because model outputs could not be explained or trusted, or because detection was never connected to a capacity to respond (Leo et al., 2019). Where these weaknesses exist, the analytics apparatus produces signals without producing safety. Conversely, institutions that manage risk well appear to do so not by owning the best single model but by integrating detection with assessment and response, which suggests that the coherence of the whole system, its detection-response integration, may matter more than any individual component (Teece, 2007).

This matters because the stakes in financial risk are asymmetric and consequential. A missed anomaly can become a large loss, a regulatory breach, or, through interconnection, a systemic event; an excess of false alarms wastes scarce investigative capacity and erodes trust in the system (Battiston et al., 2016). A detection system that is accurate but slow, or fast but opaque, or powerful but disconnected from response, fails in its essential purpose (Kanti & Rony, 2022; Sadia et al., 2022). The study therefore investigates how an AI-based risk-analytics system, understood as an integrated capability, influences financial risk-management effectiveness, and which of its components most strongly drive that effect.

Objectives of the Study

The general objective of this study was to assess how an AI-based risk-analytics system, integrating risk data aggregation, AI risk modeling, anomaly detection, early warning, model governance, and detection-response integration, influences financial risk-management effectiveness. The specific objectives were: to examine the influence of risk data aggregation and lineage on detection-response integration; to assess the relationship between AI risk-modeling capability and risk-management

effectiveness; to evaluate the effect of anomaly detection performance on effectiveness; to determine the role of early-warning and signal timeliness in strengthening effectiveness; to test the effect of detection-response integration on effectiveness; and to examine the role of model risk governance and validation in improving effectiveness. A further objective was to gauge system maturity and capability gaps across the risk-analytics lifecycle.

Research Hypotheses

Six hypotheses guided the study. H1 held that risk data aggregation and lineage significantly influence detection-response integration. H2 held that AI risk-modeling capability is significantly and positively related to financial risk-management effectiveness. H3 held that anomaly detection performance significantly influences effectiveness. H4 held that early-warning and signal timeliness significantly strengthen effectiveness. H5 held that detection-response integration significantly predicts effectiveness. H6 held that model risk governance and validation significantly strengthen effectiveness.

Significance of the Research

The study speaks to a decision that risk leaders face directly: where to invest to make risk analytics actually reduce loss. By separating the components of an AI-based risk system and estimating their contributions, it indicates whether the priority should be better data, better models, faster warnings, stronger governance, or the integration that binds detection to response (Hull, 2018). Its framing of detection and response as one system is useful to risk managers, quantitative developers, and compliance officers who must improve both together, and its maturity index and capability priority matrix turn statistical findings into an improvement agenda. For scholarship, the study connects the anomaly-detection, financial-risk, and model-governance literatures around a single capability view, and provides evidence on which parts of an AI risk system most shape the effectiveness of risk management.


LITERATURE REVIEW

The literature bearing on AI-based financial risk analytics is scattered across communities that seldom cite one another: anomaly detection in computer science, quantitative risk management in finance, model governance in regulation, and organizational capability in information systems. This review gathers them around one organizing claim, that detection and response form a single loop, and follows the path a risk signal travels: from the data that reveals it, through the models that detect it, to the governance that validates it and the response that acts on it. The review is selective rather than exhaustive, tracing the thread most relevant to the study's argument.

Risk in Complex Financial Systems

Financial risk is conventionally divided into classes, credit, market, operational, and liquidity, each with its own drivers and measurement traditions (Hull, 2018). Yet the defining feature of modern financial systems is that these classes interact and propagate: a liquidity shock can trigger credit losses, an operational failure can become a market event, and exposures link institutions into networks through which distress spreads (Battiston et al., 2016). This interconnection is why monitoring one risk class in isolation is insufficient, and why network and systemic-risk methods have gained prominence (Bianchi et al., 2021). For anomaly detection, the implication is that the events worth catching are often those that cross conventional boundaries, appearing normal within any single class but anomalous in their combination.

Figure 2: Risk Classes Monitored in Complex Financial Systems




Anomaly Detection Methods and Their Foundations

Anomaly detection is a mature field with a well-developed taxonomy distinguishing point, contextual, and collective anomalies, and a range of methods spanning statistical, distance-based, density-based, and learning-based approaches (Chandola et al., 2009). Two unsupervised methods have become especially influential in financial settings because they require no labeled events. The isolation forest scores anomalies by how easily a point is separated from the rest of the data through random partitioning, on the premise that anomalies require fewer splits to isolate (Abu Naser Md Golam & Amir, 2023; Md Tohidul, 2023).

Beyond isolation forests and density-based methods, supervised and deep-learning detectors exploit labels where available and capture complex, nonlinear structure, though at the cost of interpretability (Pang et al., 2021). These learning-based approaches have been applied widely to financial fraud and irregularity detection, from data-mining reviews of the field (Ngai et al., 2011) to machine-learning models that detect accounting fraud in public firms (Bao et al., 2020), and, in credit risk, to explainable models that pair predictive power with interpretability (Bussmann et al., 2021). The recurring lesson across this literature is that no single detector dominates; robust detection typically combines methods, which is why the study treats detection performance as a system property rather than a model property (Ahmed et al., 2016).

Figure 3: Anomaly Detection Signal Pipeline




Evaluating Detection under Extreme Imbalance

Evaluating anomaly detectors is genuinely difficult because the events of interest are rare, which makes conventional accuracy meaningless, a detector that flags nothing is right almost all of the time. Detection is therefore judged on its ability to trade the capture of true events against the burden of false alarms, using precision, recall, and their harmonic mean (Badrul & Mominul, 2023; Hossan & Adar, 2023).

These measures focus attention on the rare positive class, where precision reflects the proportion of alerts that are real and recall the proportion of real events caught (Davis & Goadrich, 2006). Under severe imbalance, precision-recall analysis is more informative than the more familiar ROC curve, because it does not let the abundant true negatives flatter the result (Saito & Rehmsmeier, 2015). The practical consequence for risk systems is that the operating threshold must be chosen deliberately,

balancing the cost of missed events against the cost of overwhelming investigators, a choice that belongs to the design of the whole system, not to the model alone.


Figure 4: Detecting True Risk without Drowning in False Alerts



Risk Data Aggregation and Lineage

Beneath every detector lies the data it consumes, and in risk analytics the quality and completeness of that data is not merely a technical concern but a regulatory one. The information-systems literature has long characterized data quality as multidimensional and use-dependent (Wang & Strong, 1996), and shown that poor data quality propagates silently and expensively through downstream analysis (Redman, 1998). In risk specifically, regulators responded to the financial crisis with principles for effective risk data aggregation and reporting, recognizing that institutions could not manage risks they could not see because their data was fragmented across systems (BCBS, 2013). The study therefore treats risk data aggregation and lineage, the ability to assemble complete, traceable, timely risk data, as a foundational capability whose weakness caps the performance of everything built above it (Einav & Levin, 2014).

Figure 5: Composing a Risk Signal from Multiple Detectors



Model Risk Governance and Explainability

As institutions rely on AI to make risk decisions, the models themselves become a source of risk, and governing that risk has become a discipline of its own. Supervisory guidance frames model risk management around the full model lifecycle, development, validation, use, and monitoring, with independent validation as a central safeguard (Federal Reserve, 2011). More recent guidance extends these expectations to complex and machine-learning models, emphasizing that opacity and drift raise, rather than lower, the bar for validation and monitoring (ECB, 2019). Governance and explainability are tightly linked: a model whose outputs cannot be explained cannot be validated, justified to regulators, or trusted by the investigators who must act on its alerts (Rudin, 2019). Attribution methods such as Shapley values have become a standard tool for restoring interpretability, allocating a model's

output across its inputs (Risha & Khalid, 2023). Related methods build local, interpretable approximations of complex models to explain individual alerts (Ribeiro et al., 2016), and the broader explainable-AI literature catalogues the techniques and their trade-offs (Barredo Arrieta et al., 2020).

Timeliness, Streaming, and Concept Drift

A risk signal has a shelf life: a warning that arrives after the loss has crystallized is of little use. This places timeliness at the center of effective detection and motivates the shift from batch to streaming analytics, in which events are scored as they arrive (Ahmad et al., 2017). Streaming settings, however, introduce concept drift, the tendency of the statistical properties of data to change over time, which degrades static models and requires detectors that adapt continuously (Gama et al., 2014). Drift is especially acute in financial risk, where behavior shifts in response to markets, regulation, and adversarial adaptation. The study captures this dimension as early-warning and signal timeliness, recognizing that when a signal arrives is as consequential as whether it arrives at all.


Theoretical Framework

The capabilities reviewed above are usually studied apart, but in an effective risk function they operate as a continuous cycle, and describing that cycle is the study's central theoretical contribution. It proposes the Observe–Detect–Assess–Respond (ODAR) sentinel loop: a situational-awareness cycle through which an AI-based risk system maintains vigilance over a complex financial environment. In the observe phase, the system aggregates and monitors risk data, drawing on data aggregation and lineage. In the detect phase, it identifies anomalies and emerging signals, exercising modeling capability and detection performance. In the assess phase, it evaluates the severity and exposure implied by a signal, invoking risk measurement and governance. In the respond phase, decision-makers act, escalate, or mitigate, and the outcomes, confirmed events, losses, false alarms, feed back as new data that recalibrates detectors and thresholds, closing the loop.

The ODAR loop rests on three theoretical foundations. From situational-awareness theory it takes its core structure: effective response depends on perceiving elements in the environment, comprehending their meaning, and projecting their implications, a progression that maps onto observe, detect, and assess (Endsley, 1995). From resilience engineering it takes the premise that safety in complex systems comes from the ongoing capacity to anticipate, monitor, respond to, and learn from disturbance, rather than from static barriers (Hollnagel et al., 2006). And from socio-technical systems theory it takes the insight that the technical detectors and the human responders must be designed together, since a detector without a response capability, or a response team without reliable signals, leaves the loop broken (Trist, 1981). The integrating construct in the framework, detection-response integration, measures how well the loop is closed, how smoothly a detected anomaly becomes an assessed and answered one (Leo et al., 2019).

Conceptual Framework

Rendering the ODAR loop as a testable model, the study positions six capabilities as drivers of financial risk-management effectiveness, with detection-response integration acting as the integrating construct through which the others largely operate. Risk data aggregation and lineage, AI risk-modeling capability, anomaly detection performance, early-warning and signal timeliness, and model risk governance and validation each contribute to effectiveness, while detection-response integration captures the coherence with which they are wired together and is expected to be the strongest single predictor (Teece, 2007). The conceptual model, shown in Figure 7, includes a feedback path from realized effectiveness back to the data and modeling capabilities, reflecting the closed-loop logic of the sentinel cycle and distinguishing it from linear views in which detection simply produces alerts (Hollnagel et al., 2006).

Figure 7: Conceptual Model for Financial Risk-Management Effectiveness

METHODS

Research Design

The study used a quantitative, cross-sectional, case-based design, suited to measuring perceptions of an AI-based risk-analytics system across a defined population at one point in time and to testing hypothesized relationships among latent constructs (Venkatesh et al., 2003). A survey strategy was appropriate because the constructs, risk data aggregation, modeling capability, detection performance, warning timeliness, governance, detection-response integration, and risk-management effectiveness, are best judged by the professionals who operate within such systems and can assess how well they function in practice. The design favors comparability across respondents over the depth of a single case, allowing the relative contribution of each capability to be estimated.

Population and Sampling

The target population comprised risk and analytics professionals engaged in financial risk assessment, anomaly detection, and model governance within complex financial data systems. A purposive sampling strategy reached respondents with direct, informed exposure to the constructs, appropriate when the aim is expert assessment rather than broad generalization (Venkatesh et al., 2003). Of 176 questionnaires distributed, 167 were returned and 158 were retained after screening for completeness and engagement, a valid response rate of 89.8%. The sample spanned risk analysts and managers, quantitative and model developers, data scientists and machine-learning engineers, compliance and audit staff, and risk-technology specialists, with 69.0% reporting direct involvement in risk modeling or anomaly detection.


Table 1: Response Rate Summary

Category	Count	Percentage
Distributed	176	100.0%
Returned	167	94.9%
Valid and used	158	89.8%
Excluded after screening	9	5.1%

Table 2: Respondent Demographics (n = 158)

Characteristic	Category	n	%
Role	Risk analysts / managers	40	25.3%
	Quant / model developers	35	22.2%
	Data scientists / ML engineers	31	19.6%
	Compliance & audit staff	28	17.7%
	Risk technology / systems	24	15.2%
Experience	1–3 years	26	16.5%
	4–6 years	66	41.8%
	More than 6 years	66	41.8%

Figure 10: Respondent Profile by Role and Experience



Instrument and Measures

Data were collected with a structured questionnaire measuring seven constructs on five-point Likert scales anchored from strongly disagree to strongly agree. Items were adapted from established instruments in anomaly detection, risk management, model governance, and information-systems research, then refined for the risk-analytics context (Leo et al., 2019). The seven constructs were risk data aggregation and lineage, AI risk-modeling capability, anomaly detection performance, early-warning and signal timeliness, model risk governance and validation, detection-response integration, and financial risk-management effectiveness. Content validity was established by expert review, and a pilot with a small group of risk professionals confirmed clarity and internal consistency before full deployment (Venkatesh et al., 2003).

DATA ANALYSIS

Analytical Approach

Analysis proceeded in stages: descriptive statistics summarized the level of each construct; Cronbach's alpha assessed internal-consistency reliability; Pearson correlations examined bivariate relationships; and multiple linear regression estimated the joint contribution of the capabilities to risk-management effectiveness. A system maturity index and a capability priority matrix then translated the statistics into practical terms. Reliability was quantified with coefficient alpha, given in Equation (1):

$$\alpha = \frac{K}{K-1} \left(1 - \frac{\sum_{i=1}^K \sigma_i^2}{\sigma_T^2} \right) \quad (1)$$

In Equation (1), K is the number of items, the numerator sums item variances, and the denominator is the variance of the total score; values at or above 0.70 indicate acceptable reliability (Leo et al., 2019). The integrating construct, detection-response integration, was constructed as a weighted mean of its component capabilities, as in Equation (2):

$$\text{DRI} = \sum_{c=1}^C w_c \bar{x}_c, \quad \sum_{c=1}^C w_c = 1, \quad w_c \geq 0 \quad (2)$$

Equation (2) aggregates the constituent capability scores into a single integration measure, with weights summing to one. The central test of the study's model used multiple regression, specified in Equation (3):

$$\text{RME} = \beta_0 + \sum_{k=1}^6 \beta_k X_k + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2) \quad (3)$$

Equation (3) regresses financial risk-management effectiveness on the six capability predictors, with standardized coefficients indicating the relative strength of each (Kou et al., 2019). The assumptions of linear regression, linearity, independence, homoscedasticity, and approximate normality of residuals, were checked before interpretation.

Descriptive Statistics

Table 3 reports the mean and standard deviation of each construct. All seven fell in the high band, indicating that respondents generally regarded their systems as capable, with meaningful variation. Financial risk-management effectiveness recorded the highest mean at 4.26, and risk data aggregation and lineage the lowest at 3.95, a pattern consistent with the study's argument that data foundations are often the weakest link even where outcomes are valued (BCBS, 2013).

Table 3: Descriptive Statistics of Study Variables

Variable	Mean	SD
Financial risk-management effectiveness	4.26	0.51
Anomaly detection performance	4.19	0.54
Detection-response integration	4.14	0.55
AI risk-modeling capability	4.10	0.58
Early-warning & signal timeliness	4.05	0.61
Model risk governance & validation	4.00	0.63
Risk data aggregation & lineage	3.95	0.66

Reliability Analysis

Table 4 presents reliability coefficients. All constructs exceeded the 0.70 threshold, with alpha ranging from 0.83 to 0.94, indicating good to excellent internal consistency and supporting the use of the scales in subsequent analysis (Leo et al., 2019). The high reliability of the effectiveness and integration scales is particularly important, since these anchor the study's central relationship.

Table 4: Reliability Analysis of Study Constructs

Construct	Items	Cronbach's Alpha	Decision
Financial risk-management effectiveness	6	0.94	Reliable
Detection-response integration	6	0.92	Reliable
Anomaly detection performance	5	0.90	Reliable
AI risk-modeling capability	6	0.88	Reliable
Early-warning & signal timeliness	5	0.86	Reliable
Model risk governance & validation	5	0.84	Reliable
Risk data aggregation & lineage	6	0.83	Reliable

Correlation Analysis

Table 5 reports Pearson correlations among the constructs. All were positive and statistically significant, consistent with an integrated system in which the capabilities move together (Mikalef et al., 2018). The strongest association was between detection-response integration and financial risk-management effectiveness, $r = 0.81$, followed by the correlation between anomaly detection performance and effectiveness, $r = 0.73$. The prominence of integration in the correlation structure foreshadows its role as the leading predictor in the regression.

Table 5: Correlation Matrix of Study Variables

Variable	Data	Model	Detect	Warn	Govern	Integ.	Effect.
Risk data aggregation	1.00						
AI risk modeling	0.61	1.00					
Anomaly detection	0.60	0.65	1.00				
Early warning & timeliness	0.58	0.60	0.63	1.00			
Model risk governance	0.57	0.58	0.57	0.59	1.00		
Detection–response integration	0.72	0.69	0.70	0.66	0.64	1.00	
Risk-management effectiveness	0.66	0.71	0.73	0.70	0.68	0.81	1.00

Figure 15: Correlation Strength among Major Study Variables

Regression Analysis

Table 6 reports the regression model summary and Table 7 the coefficients. The six capabilities jointly explained 74.1% of the variance in financial risk-management effectiveness, $R^2 = 0.741$, adjusted $R^2 = 0.729$, and the model was highly significant, $F(6,151) = 71.94$, $p < 0.001$ (Kou et al., 2019). Detection–response integration was the strongest predictor, $\beta = 0.33$, followed by anomaly detection performance, $\beta = 0.26$, and AI risk-modeling capability, $\beta = 0.22$; early warning, risk data aggregation, and model governance were significant but smaller contributors. The dominance of integration supports the study's core claim that closing the loop between detection and response, more than any single component, drives effectiveness.

Table 6: Regression Model Summary

Statistic	Value
R	0.861
R^2	0.741
Adjusted R^2	0.729
Std. error of estimate	0.272
F (6, 151)	71.94
Significance (p)	< 0.001

Table 7: Regression Coefficients

Predictor	β	t	p
Detection–response integration	0.33	4.98	< 0.001
Anomaly detection performance	0.26	3.96	< 0.001
AI risk-modeling capability	0.22	3.38	0.001
Early-warning & signal timeliness	0.19	2.94	0.004
Risk data aggregation & lineage	0.17	2.71	0.007
Model risk governance & validation	0.15	2.36	0.019

HYPOTHESIS TESTING

The six hypotheses were evaluated using the regression results together with the correlation evidence. Table 8 summarizes the outcome. All six were supported: each capability was significantly and positively associated with its hypothesized outcome, and detection–response integration's effect on effectiveness was both significant and the largest in the model, consistent with its theorized role as the integrating construct that closes the sentinel loop (Teece, 2007).

Table 8: Summary of Hypothesis Testing


Hypothesis	Path	β / r	Result
H1	Risk data aggregation → detection–response integration	r = 0.72	Supported
H2	AI risk-modeling capability → effectiveness	β = 0.22	Supported
H3	Anomaly detection performance → effectiveness	β = 0.26	Supported
H4	Early-warning & timeliness → effectiveness	β = 0.19	Supported
H5	Detection–response integration → effectiveness	β = 0.33	Supported
H6	Model risk governance → effectiveness	β = 0.15	Supported

SYSTEM MATURITY ANALYSIS

To turn the constructs into a practical diagnostic, the study computed a system maturity index across the stages of the risk-analytics lifecycle. The index is the mean of stage-level maturity scores, and the accompanying gap measure contrasts the importance and current effectiveness of each capability area. Table 9 reports maturity by stage. The overall index was 4.08, indicating generally mature systems, with data aggregation and detection tuning relatively strong and response, escalation, and governance comparatively weaker, a profile consistent with institutions that have invested in detection machinery faster than in the response and validation practices that make detection useful (Leo et al., 2019).

Table 9: System Maturity by Lifecycle Stage

Lifecycle Stage	Maturity (1–5)	Level
Data aggregation & lineage	4.16	High
Detection tuning	4.13	High
Model development	4.12	High
Monitoring & review	4.07	High
Early-warning signaling	4.05	High
Governance & validation	4.02	High
Response & escalation	3.98	Moderate–High

Figure 10: System Maturity Across Risk-Analytics Stages

RISK-ANALYTICS CAPABILITY PRIORITY MATRIX

The priority matrix places each capability area according to its importance and its current effectiveness, so that the widest gaps, areas rated important but not yet effective, rise to the top of the improvement agenda (Hull, 2018). Table 10 reports the results, ordered by gap. The largest gaps appeared in real-time anomaly scoring, model explainability, and data lineage and aggregation, indicating that respondents most wanted their systems to detect in real time, to make model outputs interpretable, and to strengthen the data foundations beneath detection. These priorities map cleanly onto the ODAR loop: they concentrate in the observe and detect phases, where the raw capacity to see and identify risk is formed (Ahmad et al., 2017).

Table 10: Risk-Analytics Capability Priority Matrix

Capability Area	Importance	Effectiveness	Gap	Priority
Real-time anomaly scoring	4.74	3.68	1.06	Very High
Model explainability	4.68	3.60	1.08	Very High
Data lineage & aggregation	4.64	3.58	1.06	Very High
Model drift monitoring	4.62	3.62	1.00	High
Systemic & network risk view	4.58	3.55	1.03	High
Alert triage & FP control	4.60	3.70	0.90	High
Response automation & escalation	4.40	3.72	0.68	Moderate
Cross-risk data integration	4.44	3.80	0.64	Moderate

FINDINGS

Assembled, the results tell a consistent story. First, respondents regard their AI-based risk-analytics systems as broadly capable: every construct sat in the high band, with effectiveness and anomaly detection performance at the top. Second, the capabilities are tightly interrelated, all correlations were positive and significant, as one expects if they are components of a single system rather than independent tools. Third, and most importantly, the capabilities jointly explain a large share of the variance in risk-management effectiveness, 74.1%, and detection-response integration is the strongest single predictor, ahead of anomaly detection performance and AI risk-modeling capability. Fourth, the maturity and priority analyses locate where the systems are strongest and weakest. Maturity is highest in data aggregation and detection tuning and lowest in response, escalation, and governance, while the widest capability gaps lie in real-time anomaly scoring, model explainability, and data lineage. The consistent implication is that the frontier of improvement is not a more powerful detector in isolation but faster, more interpretable detection built on stronger data and wired to a capable response, the work that closes the sentinel loop. This is exactly what the ODAR framework predicts: effectiveness rises when the observe, detect, assess, and respond phases reinforce one another, and the greatest returns come from strengthening the weakest link in that loop, which for these institutions lies in turning detected anomalies into timely, explainable, well-governed responses.

DISCUSSION

The results support the study's central proposition: that detection and response are one system, and that financial risk-management effectiveness depends less on any single detector than on how well the whole sentinel loop is closed. The dominance of detection-response integration as a predictor, $\beta = 0.33$,

is the clearest expression of this. Read through the Observe–Detect–Assess–Respond loop, integration measures how smoothly a detected anomaly becomes an assessed and answered one, so its strength as a predictor reflects the framework's premise that effectiveness emerges from the coherence of the whole cycle rather than from any single phase (Teece, 2007).

The strong effect of anomaly detection performance, $\beta = 0.26$, the second-largest, confirms that the quality of detection still matters greatly, integration cannot compensate for a system that fails to see. This is consistent with the anomaly-detection literature's emphasis on robust, well-evaluated detection under severe class imbalance (Chandola et al., 2009). Notably, respondents rated detection performance highly yet flagged real-time scoring and model explainability as their largest capability gaps, which suggests that the frontier is not raw detection power but timely and interpretable detection that can be trusted and acted upon (Ahmad et al., 2017). This echoes the accuracy–transparency tension in the literature: institutions want detectors that are both powerful and explainable, and the framework's treatment of these as jointly achievable through better system design, rather than as a strict trade-off, finds empirical support here (Rudin, 2019).

The significant effect of AI risk-modeling capability, $\beta = 0.22$, and the foundational role of risk data aggregation, its strong correlation with integration, $r = 0.72$, and significant regression effect, $\beta = 0.17$, together confirm that detection rests on data and models that must themselves be sound. That data aggregation recorded the lowest mean of any construct, 3.95, while remaining a significant driver and carrying one of the largest capability gaps, mirrors a persistent theme in both the data-quality and risk-data literatures: institutions cannot manage what they cannot see, and fragmented risk data quietly caps everything built upon it (BCBS, 2013).

The significant contributions of early warning, $\beta = 0.19$, and model risk governance, $\beta = 0.15$, though smaller, reinforce the framework's logic. Timeliness matters because a late signal is a useless one, consistent with the streaming and concept-drift literature (Gama et al., 2014); governance matters because an unvalidated or unexplained model cannot be trusted or defended, consistent with supervisory guidance and the explainability literature (Federal Reserve, 2011). The comparatively lower maturity of response, escalation, and governance, and the persistence of gaps there, indicate that the assessment-and-response side of the loop lags the detection side, precisely the imbalance the ODAR framework warns against, since a sentinel loop that detects but cannot respond leaves the institution aware of risks it does not address (Hollnagel et al., 2006). Taken together, the findings describe systems that are strong at observing and detecting but comparatively weaker at the timely, explainable, and well-governed work of assessing and responding, and they locate the greatest returns to improvement in closing those gaps rather than in adding detection power for its own sake (Leo et al., 2019).

CONCLUSION

This study set out to assess how an AI-based risk-analytics system, understood as an integrated capability, influences financial risk-management effectiveness, and through which components. Using a quantitative, cross-sectional design and responses from 158 risk and analytics professionals, it found that all six hypothesized relationships held and that the capabilities together explained 74.1% of the variance in effectiveness. Detection–response integration emerged as the strongest predictor, underscoring that the coherence with which detection, assessment, and response are wired together matters more than any single element. Anomaly detection performance and AI risk-modeling capability followed, while early warning, risk data aggregation, and model governance were significant but smaller contributors.

The study concludes that financial risk assessment and anomaly detection are best understood not as a pipeline that ends in an alert but as a closed sentinel loop that ends in a response and begins again with what that response reveals. The Observe–Detect–Assess–Respond cycle offers a way to see that loop whole: data is observed, anomalies detected, their severity assessed, and a response mounted whose outcomes feed back as new data. Effectiveness rises when these phases reinforce one another, and the evidence here indicates that the phases most in need of strengthening are the timely, explainable production of detections and their translation into governed, well-executed responses. For institutions, the practical message is to invest in integration, in the data foundations, real-time detection, explainability, and response capability that close the loop, rather than in isolated detectors; for scholarship, the study shows the value of joining the anomaly-detection, risk-management, and model-

governance literatures around a single capability view of financial risk.

RECOMMENDATIONS

Several recommendations follow from the findings. First, institutions should treat risk data aggregation and lineage as a first-order investment, strengthening the completeness, traceability, and timeliness of risk data, since data aggregation was both the lowest-rated capability and a significant driver of effectiveness, and its weakness caps detection. Second, they should move anomaly detection toward real-time, drift-aware scoring, which respondents identified as their most important capability gap, so that risks are caught while there is still time to act.

Third, institutions should invest in model explainability, adopting attribution and interpretable-modeling methods so that alerts can be understood, trusted, validated, and defended to regulators. Fourth, they should strengthen the response side of the loop, alert triage, false-positive control, escalation, and response automation, so that detected anomalies reliably become timely actions. Fifth, they should mature model risk governance and validation in step with model complexity, ensuring that opaque and adaptive models are validated and monitored rather than merely deployed. Finally, institutions should manage AI-based risk analytics as an integrated sentinel system, using a maturity assessment and a capability priority matrix to direct improvement toward the weakest phases of the Observe–Detect–Assess–Respond loop rather than toward isolated detectors.

LIMITATIONS OF THE STUDY

This study has several limitations. First, the cross-sectional design captured perceptions at a single point in time and cannot establish causal direction or track how capabilities and effectiveness co-evolve; a longitudinal design would allow stronger causal inference and better capture the closed-loop dynamics central to the ODAR framework. Second, the study relied on perceptual, self-reported measures completed by professionals, so the results reflect informed assessment rather than objective measurement of detection accuracy or loss avoidance; future work could pair perception data with hard metrics such as measured detection rates, false-positive rates, alert-to-action times, and realized losses. Third, the purposive sample, while appropriate for reaching knowledgeable respondents, limits generalizability to the broader population of institutions. Fourth, the constructs were assessed at a general level, and the relative importance of specific capabilities may vary across banking, insurance, capital markets, and payments, and across regulatory regimes and firm sizes. Finally, the study did not directly measure the cost or organizational-change effort required to raise system maturity, which is an important practical consideration and a promising avenue for future research.

REFERENCES

- [1]. Abu Naser Md Golam, M., & Amir, R. (2023). Digital Twin Modeling of Fiber Optic Telecom Infrastructure for Proactive SCADA Network Resilience in Power Transmission Utilities. *International Journal of Scientific Interdisciplinary Research*, 4(4), 413–448. <https://doi.org/10.63125/04mgf566>
- [2]. Amir, R., & Chapal, B. (2022). Federated Learning Architectures for Distributed Engineering Project Knowledge Management A Simulation-Based Analysis. *Review of Applied Science and Technology*, 1(04), 411–446. <https://doi.org/10.63125/4k4pjr64>
- [3]. Kazi Rakib Hasan, S., & Uddin, H. M. M. (2022). Scalable AI For Project Portfolio Management: A Mixed-Methods Study Combining Distributed Computing Benchmarks. *Review of Applied Science and Technology*, 1(04), 375–410. <https://doi.org/10.63125/0kk4wf20>
- [4]. Md Tohidul, I. (2023). ERP-Integrated Financial Analytics Systems for Corporate Financial Reporting, Cash Flow Analysis, and Strategic Investment Planning. *American Journal of Advanced Technology and Engineering Solutions*, 3(04), 87–128. <https://doi.org/10.63125/qhk96h35>
- [5]. Md. Abdur, R., & Iftekhar, A. (2021). Customer Retention Forecasting in Mobile Wallet Services Using Neural Networks: A Comparative Quantitative Study. *International Journal of Business and Economics Insights*, 1(4), 70–102. <https://doi.org/10.63125/dyrpc387>
- [6]. Mohammad Badrul, A., & Md. Mominul, H. (2023). Machine Learning-Based Assessment of Environmental and Public Health Risks in Sewerage Sludge Management Systems. *American Journal of Advanced Technology and Engineering Solutions*, 3(02), 33–77. <https://doi.org/10.63125/y3yfxa92>
- [7]. Mohammad Shoriful Hossan, S., & Adar, C. (2023). SIL-Compliant Safety Instrumented System Design for Oil & Gas Hazardous-Area Operations Using IEC 61508 / IEC 61511 Frameworks. *Journal of Sustainable Development and Policy*, 2(02), 43–85. <https://doi.org/10.63125/cb3d3w81>
- [8]. Muhammad Mohiul, I., & Mohammad Badrul, A. (2022). Predictive Analytics for Optimizing Sewerage Sludge Treatment Efficiency and Resource Recovery in Dhaka City. *Journal of Sustainable Development and Policy*, 1(04), 158–200. <https://doi.org/10.63125/t9zdp458>
- [9]. Mukut Kanti, B., & Rony, S. (2022). Quantitative Risk Assessment and Safety Performance Benchmarking for Structural Integrity Management in Retail Supply Chain Infrastructure. *American Journal of Data Science and*

- Analytics, 3(12), 42-88. <https://doi.org/10.63125/n88vth90>
- [10]. Risha, A., & Kazi Mohammad Khalid, A. (2023). A Meta-Analysis of AI-Driven Geospatial Analytics for Predictive Maintenance of Critical Infrastructure in Developing Economies. *International Journal of Scientific Interdisciplinary Research*, 4(4), 375–412. <https://doi.org/10.63125/rayrex49>
- [11]. Sadia, Z., Mohammad Ariful, I., Sumyta, H., & Hosne Ara, M. (2022). A Review of AI-Powered Data Visualization in Enterprise Reporting: Dashboard Design And Interactive Analytics. *American Journal of Advanced Technology and Engineering Solutions*, 2(01), 32-54. <https://doi.org/10.63125/gabst658>
- [12]. Ahmad, S., Lavin, A., Purdy, S., & Agha, Z. (2017). Unsupervised real-time anomaly detection for streaming data. *Neurocomputing*, 262, 134–147. <https://doi.org/10.1016/j.neucom.2017.04.070>
- [13]. Ahmed, M., Mahmood, A. N., & Islam, M. R. (2016). A survey of anomaly detection techniques in financial domain. *Future Generation Computer Systems*, 55, 278–288. <https://doi.org/10.1016/j.fjncs.2015.11.016>
- [14]. Bao, Y., Ke, B., Li, B., Yu, Y. J., & Zhang, J. (2020). Detecting accounting fraud in publicly traded U.S. firms using a machine learning approach. *Journal of Accounting Research*, 58(1), 199–235. <https://doi.org/10.1111/1475-679X.12292>
- [15]. Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., ... Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- [16]. Basel Committee on Banking Supervision. (2013). Principles for effective risk data aggregation and risk reporting (BCBS 239). Bank for International Settlements. <https://www.bis.org/publ/bcbs239.htm>
- [17]. Battiston, S., Farmer, J. D., Flache, A., Garlaschelli, D., Haldane, A. G., Heesterbeek, H., ... Scheffer, M. (2016). Complexity theory and financial regulation. *Science*, 351(6275), 818–819. <https://doi.org/10.1126/science.aad0299>
- [18]. Bianchi, D., Büchner, M., & Tamoni, A. (2021). Bond risk premia with machine learning. *The Review of Financial Studies*, 34(2), 1046–1089. <https://doi.org/10.1093/rfs/hhaa062>
- [19]. Breunig, M. M., Kriegel, H.-P., Ng, R. T., & Sander, J. (2000). LOF: Identifying density-based local outliers. *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, 93–104. <https://doi.org/10.1145/342009.335388>
- [20]. Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable machine learning in credit risk management. *Computational Economics*, 57(1), 203–216. <https://doi.org/10.1007/s10614-020-10042-0>
- [21]. Cao, L. (2022). AI in finance: Challenges, techniques, and opportunities. *ACM Computing Surveys*, 55(3), 1–38. <https://doi.org/10.1145/3502289>
- [22]. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 1–58. <https://doi.org/10.1145/1541880.1541882>
- [23]. Davis, J., & Goadrich, M. (2006). The relationship between precision-recall and ROC curves. *Proceedings of the 23rd International Conference on Machine Learning*, 233–240. <https://doi.org/10.1145/1143844.1143874>
- [24]. Dixon, M. F., Halperin, I., & Bilokon, P. (2020). *Machine learning in finance: From theory to practice*. Springer. <https://doi.org/10.1007/978-3-030-41068-1>
- [25]. Einav, L., & Levin, J. (2014). Economics in the age of big data. *Science*, 346(6210), 1243089. <https://doi.org/10.1126/science.1243089>
- [26]. Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32–64. <https://doi.org/10.1518/001872095779049543>
- [27]. European Central Bank. (2019). ECB guide to internal models: Risk-type-specific chapters. European Central Bank. https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.trim_guide_risk_type_specific_chapters_201910.en.pdf
- [28]. Federal Reserve & Office of the Comptroller of the Currency. (2011). Supervisory guidance on model risk management (SR 11-7). Board of Governors of the Federal Reserve System. <https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm>
- [29]. 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>
- [30]. Goodell, J. W., Kumar, S., Lim, W. M., & Pattnaik, D. (2021). Artificial intelligence and machine learning in finance: Identifying foundations, themes, and research clusters from bibliometric analysis. *Journal of Behavioral and Experimental Finance*, 32, 100577. <https://doi.org/10.1016/j.jbef.2021.100577>
- [31]. Hollnagel, E., Woods, D. D., & Leveson, N. (Eds.). (2006). *Resilience engineering: Concepts and precepts*. Ashgate. <https://doi.org/10.1201/9781315605685>
- [32]. Hull, J. C. (2018). *Risk management and financial institutions* (5th ed.). Wiley. <https://www.pearson.com/en-us/subject-catalog/p/risk-management-and-financial-institutions/P200000005938>
- [33]. Jorion, P. (2007). *Value at risk: The new benchmark for managing financial risk* (3rd ed.). McGraw-Hill. <https://www.mheducation.com/highered/product/value-risk-3rd-edition-jorion/M9780071464956.html>
- [34]. Kou, G., Chao, X., Peng, Y., Alsaadi, F. E., & Herrera-Viedma, E. (2019). Machine learning methods for systemic risk analysis in financial sectors. *Information Sciences*, 470, 111–129. <https://doi.org/10.1016/j.ins.2018.10.023>
- [35]. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- [36]. Leo, M., Sharma, S., & Maddulety, K. (2019). Machine learning in banking risk management: A literature review. *Risks*, 7(1), 29. <https://doi.org/10.3390/risks7010029>
- [37]. 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>

- [38]. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765–4774. <https://doi.org/10.48550/arXiv.1705.07874>
- [39]. McNeil, A. J., Frey, R., & Embrechts, P. (2015). *Quantitative risk management: Concepts, techniques and tools* (Rev. ed.). Princeton University Press. <https://press.princeton.edu/books/hardcover/9780691166278/quantitative-risk-management>
- [40]. Mikalef, P., Pappas, I. O., Krogstie, J., & Giannakos, M. (2018). Big data analytics capabilities: A systematic literature review and research agenda. *Information Systems and e-Business Management*, 16(3), 547–578. <https://doi.org/10.1007/s10796-017-9720-z>
- [41]. Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review. *Decision Support Systems*, 50(3), 559–569. <https://doi.org/10.1016/j.dss.2010.08.006>
- [42]. Ozbayoglu, A. M., Gudelek, M. U., & Sezer, O. B. (2020). Deep learning for financial applications: A survey. *Applied Soft Computing*, 93, 106384. <https://doi.org/10.1016/j.asoc.2020.106384>
- [43]. Pang, G., Shen, C., Cao, L., & Van Den Hengel, A. (2021). Deep learning for anomaly detection: A review. *ACM Computing Surveys*, 54(2), 1–38. <https://doi.org/10.1145/3439950>
- [44]. Redman, T. C. (1998). The impact of poor data quality on the typical enterprise. *Communications of the ACM*, 41(2), 79–82. <https://doi.org/10.1145/269012.269025>
- [45]. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”: Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144. <https://doi.org/10.1145/2939672.2939778>
- [46]. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- [47]. Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. *PLOS ONE*, 10(3), e0118432. <https://doi.org/10.1371/journal.pone.0118432>
- [48]. Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal*, 28(13), 1319–1350. <https://doi.org/10.1002/smj.640>
- [49]. Trist, E. (1981). *The evolution of socio-technical systems: A conceptual framework and an action research program*. Ontario Quality of Working Life Centre. <https://doi.org/10.4324/9781315227979>
- [50]. Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>
- [51]. Wang, R. Y., & Strong, D. M. (1996). Beyond accuracy: What data quality means to data consumers. *Journal of Management Information Systems*, 12(4), 5–33. <https://doi.org/10.1080/07421222.1996.11518099>
- [52]. West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. *Computers & Security*, 57, 47–66. <https://doi.org/10.1016/j.cose.2015.09.005>