

AI-Driven Governance, Risk and Compliance (GRC) for Financial Markets

 **Abhishek Kakkar**¹

¹ Independent Researcher, USA

Email: achievementgrow@gmail.com

 **Dr. Sonal Kapoor**²

² Associate Professor

Email: sonalkapoor.rk@its.edu.in

ABSTRACT: The rapid expansion of Financial Technology (FinTech) platforms has fundamentally transformed different areas of delivery of financial services with business expansion into real-time payments, algorithmic trading, digital lending, decentralized financial products and so on. Today financial institutions face a compliance environment that has outgrown the tools built to manage it. Regulatory obligations increase and multiply across jurisdictions of operation; financial crime volumes continue to rise with advancement of technology, and data generation has grown exponentially to a scale where manual oversight is structurally impossible at the transaction level.

Legacy Governance, Risk and Compliance (GRC) frameworks were built around periodic audit cycles, static rule sets, and report heavy workflows designed for a different era. GRC frameworks are not failing because organizations run them poorly; they are failing because the environment has changed faster than the frameworks evolved. Artificial intelligence is the mechanism through which that gap is being closed in real-time, and this paper examines what that looks like in practice.

This paper covers what actually happens when financial institutions deploy AI in their GRC functions. Where it works, where it falls short, and what separates the two. The analysis draws data based on three peer-reviewed empirical studies covering over 1,155 publications and survey data from 564 financial sector professionals. Those studies are cited as evidence, not as subject matter. Some of the financial market scenarios are examined through the lens of what AI has demonstrably changed.

There are consistent findings those are worth stated across each domain from AML surveillance and KYC automation in retail banking, Credit underwriting in commercial lending, Regulatory reporting under Basel IV, and Board transparency in listed corporates.

AI returns are not determined by the sophistication of the algorithm. They are determined by the quality of the governance infrastructure surrounding it. Institutions that have learned this art early on are compounding the advantage. Those that have not are finding that AI investments disappoint not because the technology fails, but because the data and oversight conditions necessary for it to succeed were never put in place.

AI's contribution to GRC is real, statistically validated, and growing but it is not unconditional. Institutions that invest in data quality, lineage, metadata, model oversight, and ethical AI frameworks extract substantially stronger GRC returns from AI than those that deploy AI tools without the underlying governance infrastructure. This is not a theoretical nuance; it is the difference between an AI system that catches financial crime and one that floods investigators with noise.

KEYWORDS: Artificial Intelligence (AI), Governance, Risk and Compliance (GRC), Financial Markets, FinTech, Machine Learning, Regulatory Compliance, Data Governance, Anti-Money Laundering (AML), Explainable Artificial Intelligence (XAI), Model Risk Management, Financial Crime Detection.

1. Introduction

1.1. The FinTech Revolution and the GRC Challenge

Two decades ago, bank's compliance function dealt with a manageable universe of products, counter parties, and transactions. Today it deals with open banking APIs feeding dozens of third-party channels, instant payment rails that settle in milliseconds, digital lending decisions made algorithmically at volume. The customer expects frictionless onboarding regardless of the underlying verification and validation requirements. The technology that made those capabilities possible did not come with a governance manual. The Basel Committee has progressively tightened capital adequacy, liquidity, and stress testing requirements through Basel III and IV, each iteration demanding greater precision in risk measurement and data governance than the last. What has not kept pace is the tools available to compliance and risk teams to actually satisfy those demands at the speed and scale modern finance grows. The question is no longer whether artificial intelligence (AI) belongs in the GRC function; it is how to deploy it effectively and at scale.

A retail bank in 2026 can onboard a customer in five minutes on a mobile app, execute a cross-border payment in seconds, and extend a credit line based on a real-time ML scoring decision; all before a human compliance officer has seen any part of the interaction. That is the commercial reality driving customer acquisition and revenue growth. It is also a compliance architecture that would have been unrecognizable to the risk managers who built the GRC frameworks currently governing it.

At the same time, regulatory obligations have multiplied. In the United States, the Dodd-Frank Act and SEC reporting requirements create layers of compliance obligations. The European Union's Markets in Financial Instruments Directive (MiFID II), General Data Protection Regulation (GDPR), and the Digital Operational Resilience Act (DORA) impose stringent requirements on data handling, algorithmic transparency, and operational risk management. The Financial Action Task Force (FATF) continues to raise AML and counter-terrorism financing standards globally. FATF estimates between **2% and 5%** of global GDP as much as **USD 2 trillion** annually is laundered through the financial system. Global financial institutions paid over **USD 320 billion** in regulatory fines between 2008 and 2023, a significant proportion related to

failures in AML controls, sanctions compliance, and financial reporting (BCG, 2023). Financial institutions now routinely operate under simultaneous obligations to dozens of regulatory bodies across multiple jurisdictions.

GRC functions have historically been the last area to receive meaningful technology investment, and the consequences are visible. Transaction volumes, regulatory texts, customer records, and reporting obligations have all grown faster than the headcount and tooling available to manage them. Choowan and Daovisan (2026) documented in their systematic review of 1,155 peer-reviewed studies spanning a decade of financial services research, that data silos, legacy system fragmentation, and unclear lineage remain persistent features of the institutional landscape. These are not cosmetic problems. They undermine the reliability of the compliance and risk processes built on top of them, and it is against this backdrop that AI has entered the GRC function.

1.2 What AI Actually Does in Financial GRC

The term 'AI' is used broadly, so it is worth being specific about what it means in the GRC context. In financial services, AI-driven GRC draws primarily on three capability families.

- **Machine learning (ML)** models detect patterns in large datasets, identifying anomalies, predicting defaults, or classifying transactions with a statistical sophistication that rule-based systems cannot replicate.
- **Natural Language Processing (NLP)** enables machines to read, interpret, and act on unstructured text, regulatory documents, customer emails, loan agreements, news feeds at a scale and speed no human team can match.
- **Deep learning (DL)** and **hybrid AI** architectures combine multiple model types to address complex, multi-dimensional risks across diverse data sources.

Each of these are worth examining as a deployed reality, not as a proof-of-concept.

1. **JP Morgan Chase's** Contract Intelligence (COiN) platform, documented by Artene, Domil and Ivascu (2024), processes commercial credit agreements using NLP to extract critical data points and clauses; work

that previously consumed 360,000 hours of lawyer and analyst time annually.

2. **Deloitte's ACL Analytics** platform applies ML to financial audit data to surface anomalies and patterns warranting scrutiny, enabling auditors to concentrate their judgment on high-risk areas rather than routine data review.
3. **Al-Wasleh**, a Jordanian leasing firm examined by Shaban and Omoush (2025), deployed ML-based credit scoring that analyzes transaction history, payment behavior, and socio-economic indicators in real time to generate credit assessments, supplemented by an AI-integrated ERP system automating accounts reconciliation, anomaly detection, and regulatory compliance checks.

The common thread across these implementations is not the sophistication of the algorithm, it is the measurable reduction in risk, cost, and error that AI delivers when applied to tasks that are high-volume, pattern-dependent, and time-sensitive. These are precisely the characteristics of GRC work: monitoring millions of transactions for suspicious patterns, reading thousands of pages of regulatory updates, reviewing hundreds of loan applications, reconciling vast datasets for regulatory reporting.

1.3 Where Are the Research Gaps

There is substantial literature available on AI in financial services. What it has been slower to address is a set of questions that practitioners encounter every implementation cycle and that the academic record has not yet resolved cleanly.

1. The relationship between AI capability and GRC outcome is frequently presented as direct and automatic '**deploy AI, improve compliance**'. Whereas the empirical evidence tells a more nuanced story. Choowan and Daovisan (2026) confirm through structural equation modelling (SEM) that governance maturity mediates the AI-to-outcome relationship with a coefficient of 0.73 ($p < 0.001$). The quality of the data governance framework surrounding an AI deployment is as determinative of outcome as the algorithm itself. This finding is under represented in practitioner discourse.
2. Critical practical gap is **explainability**, AI-driven

governance models frequently struggle with explainability, making it difficult for compliance officers, internal auditors, and external regulators to understand AI-generated decisions. This limitation is acute in high-stakes contexts, credit decisions, AML alert adjudication, and model risk management; where regulatory expectations demand that financial institutions be able to explain and defend their model outputs. Artene et al. (2024) notes that even a three-layer neural network modest by modern standards presents significant interpretability challenges when deployed in financial decision-making environments where transparency and accountability are legally mandated. UK Financial Conduct Authority (FCA) and US Office of the Comptroller of the Currency (OCC) model risk management guidelines both require institutions to be able to explain and defend model outputs. Institutions deploying AI in consequential decision contexts without a credible explainability strategy are accumulating regulatory exposure that is not yet fully materialized.

3. There is a notable absence of research on **cross-jurisdictional AI-GRC integration**; how firms operating across multiple regulatory regimes can build unified AI-governance architectures that satisfy divergent requirements simultaneously. The interplay between GDPR's data minimization requirements and the data intensity demands of modern ML models remains unresolved in both literature and regulatory guidance.

1.4 Research Objectives and Hypotheses

This paper sets out to answer four questions that matter to practitioners, not just for academics:

- How is AI currently changing the way financial institutions manage regulatory compliance, and what determines whether those changes are effective?
- Where in the GRC landscape (AML, credit risk, market risk, regulatory reporting, audit) is AI generating the most significant and measurable improvement.
- What does the empirical evidence tell us about the governance maturity, data quality, model risk oversight that institutions must build before AI can reliably deliver GRC value?

- What should financial institutions do now, and what should regulators require, to ensure that AI-driven GRC serves the interests of financial stability and consumer protection?

The paper tests four hypotheses, each grounded in and validated against the empirical evidence of the three supporting studies:

1. AI integration into GRC functions significantly improves regulatory compliance effectiveness, with governance maturity acting as a mediating condition that determines the magnitude of improvement.
2. AI-driven risk analytics encompassing credit scoring, fraud detection, and market risk modelling. Deliver materially superior outcomes to rule-based predecessors, conditional on the quality of underlying data governance.
3. AI-enabled financial transparency strengthens corporate governance structures including board-level decision quality, stakeholder engagement, and internal audit effectiveness.
4. The financial value generated by AI-driven GRC is not uniformly distributed; institutions with advanced governance maturity capture materially stronger returns than those at early implementation stages.

2. Methodology

2.1 Research Design

This study uses mixed-methods design. The quantitative backbone drawn on published findings from three peer-reviewed studies, allowing hypothesis testing against robust, independently validated datasets. The qualitative dimension incorporates real-world implementation use cases drawn from both the supporting literature and industry reporting, providing the practitioner context that purely statistical analysis cannot supply.

A mixed-methods approach is appropriate here because the research span multiple levels of analysis. Research helps understand whether AI improves AML detection rates requires quantitative evidence. Understand why a particular implementation success or failure requires qualitative examination of the institutional conditions, governance structures, data quality, staff capability, leadership commitment in which it occurred. These three supporting studies span both dimensions, providing quantitative structural models alongside case evidence from real institutions. Each paper contributes something the others cannot, and the conclusions this paper draws depend on all three.

2.2 Data Sources

Table 1: Data Sources

Study	Method	Sample	Contribution to This Paper
Choowan & Daovisan, 2026	PRISMA systematic review + SEM + Quantile Radar Regression	1,155 Scopus-indexed studies (2015–2025); $\kappa = 0.81$	Structural validation of AI-governance-outcome relationship across financial sectors
Artene, Domil & Ivascu, 2024	Foundational research; neural network modelling; game theory payoff analysis; industry use case	Theoretical framework + case studies (JP Morgan, Deloitte)	Technical architecture evidence; AI decision-support modelling; industry deployment cases
Shaban & Omoush, 2025	Stratified survey; SEM (PLS-SEM); multiple regression	564 corporate professionals across banking, insurance and services sectors	Primary survey validation of AI governance impact; corporate transparency and board-level evidence
Supplementary industry sources	Industry reports, regulatory guidance	Global financial markets	Secondary analysis to contextualize and validate findings

The systematic review conducted by Choowan and Daovisan (2026) using PRISMA 2020 guidelines screened 49,037 records from Scopus database. After deduplication, abstract screening, and full-text eligibility review, 1,155 studies were retained. Study quality was assessed using the Mixed Methods Appraisal Tool (MMAT), with 48.2% rated high quality ratings.

The primary survey data of Shaban and Omoush (2025) was collected using stratified random sampling across 564 corporate governance professionals in Jordan, encompassing banking (24.8%), services (31.2%), and other financial sectors with sample size validated against Cohen's power analysis and Krejcie and Morgan's sampling formula.

Both studies applied multiple independent sensitivity and robustness tests to confirm that their findings are stable and not artifacts of sampling or specification choices.

2.3 Selection Criteria and Justification

These papers were selected for inclusion as supporting evidence on three criteria:

- a. Empirical data and analysis (not conceptual-only)
- b. Direct relevance to AI in financial GRC (practitioner) and
- c. Methodological analysis to support hypothesis testing

The three selected studies together cover a ten-year literature synthesis, primary survey data from active practitioners, and real-world implementation use cases; providing triangulated evidence across research method types.

The mixed-methods design directly addresses the core research questions. The quantitative SEM results validate the structural hypotheses about AI's effect on governance outcomes. The qualitative case evidence, COiN at JP Morgan, ACL Analytics at Deloitte, credit scoring and ERP automation at Al-Wasleh all validate that the statistical relationships observed translate into operational reality in actual financial institutions. Together they support conclusions that none of the methods alone could sustain.

3. Analytical Approach

3.1 Structural Equation Modelling – The Governance Mediation Framework

The analytical backbone of this study is the governance framework validated by Choowan and Daovisan (2026) through Structural Equation Modelling (SEM). Their three-construct model covers AI Integration, Data Governance Maturity, and Financial Decision Outcomes. They achieved excellent global fit: CFI = 0.947, RMSEA = 0.041, SRMR = 0.038, chi-square to degrees-of-freedom ratio of 2.14. These values comfortably satisfy published SEM adequacy thresholds providing a statistically defensible structural model.

Three structural paths are particularly instructive from a GRC practitioner perspective. They carry the analytical weight, and taken together they reframe how AI investment decisions should be made.

1. The path from AI Integration to Data Governance Maturity ($\beta = 0.76$, $t = 13.42$, $p < 0.001$) shows that deploying AI technologies strengthens governance by automating data lineage, enriching metadata, tightening compliance monitoring controls.
2. The path from Data Governance Maturity to Financial Outcomes ($\beta = 0.73$, $p < 0.001$) confirms that governance maturity and quality is not merely a regulatory checkbox, but an operational driver of decision making. Algorithm sophistication helps make genuine and quality decisions.
3. The direct AI-to-Outcomes path ($\beta = 0.71$) indicates with partial mediation, AI delivers both direct benefits and governance-mediated benefits, but the bulk of its value flows through the governance infrastructure it enables and requires.

For financial institutions, these finding reframes the investment question. The institutions ask '**which AI vendor should we select?**', are asking the wrong question first. The first question to be asked is: '**Is our data governance infrastructure mature enough to extract value from AI?**' The SEM evidence says the answer to the second question determines the answer to the first.

3.2 Predictive Weight Analysis – Where AI Delivers Most

Multi-Supported Vector Regression (MSVR) analysis performed by Choowan and Daovisan (2026) assigns

predictive weights to AI techniques and governance dimensions across financial decision outcomes. These weights provide a data-driven prioritization framework for GRC investment decision.

Table 2: MSVR Predictive Weights and Financial Market Interpretation (Adapted from Choowan & Daovisan)

AI Technique / Governance Dimension	MSVR Weight	What This Means in GRC Practice
Machine Learning (ML)	0.87	Highest-performing technique across AML transaction monitoring, credit risk scoring, and fraud detection. Dominant where pattern recognition across large datasets is the task.
Regulatory Compliance (Governance)	0.84	The single strongest governance predictor of outcomes. Compliance automation, regulatory change monitoring, reporting, surveillance. Delivers outsized returns at advanced maturity levels.
Risk Analytics (Governance)	0.79	Strong across banking and capital markets use cases. VaR modelling, stress testing, credit portfolio monitoring, early-warning systems.
Hybrid AI Models	0.76	Integrated architectures combining supervised ML, rules engines and NLP. Strongest at advanced governance maturity ($Q0.75 \beta = 0.83$). Best suited for complex AML-KYC platforms and integrated risk frameworks.
Natural Language Processing (NLP)	0.74	Regulatory text parsing, KYC document extraction, SAR narrative generation, contract review. The backbone of compliance automation for document-intensive processes.
Fraud Detection (Outcome)	0.81	Consistently highest-weight outcome. Payment fraud, account takeover, trade surveillance, insurance fraud. Directly monetizable through false-positive reduction and loss prevention.
Credit Scoring (Outcome)	0.78	Consumer, SME and corporate lending. Measurable improvement in Probability of Default accuracy, faster underwriting, reduced adverse selection.
Portfolio Risk Management (Outcome)	0.73	Asset allocation optimization, hedge effectiveness monitoring, macro scenario generation for ICAAP and ILAAP stress tests.

3.3 Quantile Radar Analysis – The Maturity Gradient

One of the most practically important findings from

Choowan and Daovisan (2026) is the quantile radar regression, which shows how AI-governance returns evolve with institutional maturity. This is not a single-coefficient finding, it is a gradient:




Fig 1: Quantile Regression — AI Performance by Governance Maturity (Adapted from Choowan & Daovisan)

The gradient is steeper than it first appears. Compliance automation at advanced maturity ($\beta = 0.86$) outperforms the same capability at nascent maturity ($\beta = 0.63$) by 37%. For fraud detection, the differential is 42%. The governance dividend is compounding returns on governance infrastructure investment that advanced institutions are capturing and early-stage institutions are not. The practical implication is that a mid-size bank deciding whether to invest in data catalog tools or a new AML AI vendor should invest in the data catalog first. The AI vendor's product will perform better as a result.

3.4 Regression Analysis Validation – Corporate Governance Dimension

The regression analysis of Shaban and Omoush (2025), conducted survey across 564 financial sector professionals, provides independent validation of AI's governance effects at the institutional level. Simple regression models confirm statistically strong and significant relationships between AI adoption and each governance dimension tested:




Fig 2: AI Impact on Corporate Governance Dimensions — Regression Results (Shaban & Omoush)

4. Validation Strategy

4.1 Statistical Robustness

Before drawing conclusions from the structural model, it is worth being explicit about what the underlying studies did to test whether their findings hold. Choowan

and Daovisan (2026) ran four separate sensitivity procedures on their SEM, each designed to probe a different potential weakness.

1. Exclusion of low-quality studies rated by MMAT produced negligible model fit change, confirms that findings are not driven by weaker evidence.

2. Jackknife resampling with 5% perturbation produced ($\beta < 0.02$) across all structural paths
3. Quantile invariance testing returned $q = 0.91$ ($p < 0.001$) and
4. Multi-group chi-square testing confirmed latent construct invariance, meaning the constructs measure the same things across different institutional contexts.

Collectively, these tests establish that the structural relationships identified are robust, stable, and not artifacts of sample composition or model specification.

The 5000-iteration bootstrap applied by Shaban and Omoush (2025) to their PLS-SEM paths confirmed coefficient stability, with outer loadings consistently exceeding 0.70 threshold across all constructs.

Temporal stability is important for a field evolving as rapidly as AI in financial services. Choowan and Daovisan (2026) address this directly through their multi-group SEM comparison of the early adoption period (2015-19) against the governance-mature period (2020-25). While direct path coefficients strengthened over time confirming AI's increasing operational embedding; the indirect mediation effect remained stable across both periods. This temporal invariance is significant: it means the governance-mediation relationship is not a transitional artefact of early AI deployment but a durable structural feature of how AI creates GRC value in financial institutions.

4.2 Cross-Platform Validation

Shaban and Omoush (2025) conducted parallel analysis on two independent analytical platforms SPSS V23 for traditional regression and SmartPLS for PLS-SEM and found consistent results across both. Path coefficients from SmartPLS aligned with coefficients from SPSS regression.

- a) AI-Driven $\rightarrow$ Stakeholder Engagement = 0.967
- b) AI-Driven $\rightarrow$ Executive Compensation = 0.915

The alignment between platforms supports the interpretation that these are genuine structural relationships rather than estimation artefacts, which is particularly important given the high R^2 values that the study reports – values that, without validation, might

reasonably prompt questions about model overfit.

5. Relevance to Previous Publications

5.1 Theoretical Foundations

Two theoretical frameworks provide the conceptual ground for the empirical evidence reviewed here, and it is worth being explicit about what each one explains as they explain different things.

1. **Agency Theory**, as articulated by Jensen and Meckling (1976), starts from the observation that the people who manage financial institutions and the people who own or regulate them have different information. That asymmetry is the root cause of most governance problems, boards that cannot see what management sees cannot effectively oversee it; regulators that cannot reconstruct decision logic cannot effectively challenge it. AI-driven transparency directly reduces the asymmetry through real-time reporting, automated audit trails, and AI-powered disclosure mechanisms. In terms of Agency Theory, R^2 of 0.582 on board level decision quality that Shaban and Omoush (2025) report is an empirical support to the quantitative measure of agency cost reduction in a financial services context.
2. **Institutional Theory** (DiMaggio & Powell, 1983; Scott, 2014) explains a different question: why AI adoption in financial GRC has accelerated regardless of implementation & data quality. Financial institutions are not adopting AI purely because they have done rigorous cost-benefit analyses. Institutions are adopting it partly because regulators are expecting it, peers are using it, boards are asking it, and the legitimate cost of not having AI is rising. It explains the pace of adoption, and it suggests governance gaps need to be addressed.

5.2 Extension of Prior AI-Finance Literature

The MSVR and quantile regression methodology introduced by Choowan and Daovisan (2026) extends prior AI-finance research by capturing what standard regression cannot; the non-linear, maturity-conditioned nature of AI-governance returns. Earlier studies including Avelar and Jordao (2025) and Gyau (2024) on AI's contribution gave positive financial effects without disaggregating by governance maturity.

Relative to the Jordan-focused governance literature, Shaban and Omoush (2025) extend Omoteso and Mobolaji's (2020) findings on AI's role in boardroom transparency and work on AI and financial reporting. Their contribution is the demonstration that AI governance is not confined to developed economies with mature regulatory frameworks but are observable and more impactful in the developing economy where traditional governance mechanisms are weaker.

Artene et al. (2024) contribute an original integration of game theory and neural networks in the 'Decision Harbor AI' architecture combining a three-layer neural network (31 inputs, 16 hidden units, 2 outputs). Their payoff analysis is instructive for GRC investment decisions: when both management and shareholders adopt AI-driven financial reporting, the outcome is High/High. When one party adopts and the other does not, payoffs are asymmetric.

6. Discussion and Analysis

6.1 AI in AML and Financial Crime Surveillance

The Scale Problem That Rules Cannot Solve

A global bank processing 50 million transactions per day cannot apply meaningful human judgment to more than a fraction of them. Traditional rules-based AML systems attempt to solve this by defining controls about fixed thresholds, flag any cash transaction above \$10,000, any wire transfer to a high-risk jurisdiction, any account with unusual velocity. The result is a system that is simultaneously under-inclusive (sophisticated criminal networks design around known rules) and massively over-inclusive (false positive rates of 95–98% are common in large institution AML programs). Compliance teams spend the bulk of their capacity reviewing alerts that turn out to be legitimate, rather than investigating genuinely suspicious activity.

ML-based transaction monitoring changes the architecture of this problem. Rather than applying static thresholds to individual transactions, ML models learn the behavioral signature of a customer typical transaction volumes, counterparties, timing, geography and flag deviations from that baseline. Network graph analysis maps relationships between accounts to surface layering structures that only become visible across the full picture of flows. NLP processes unstructured news, sanctions list, and adverse media at a speed and coverage no human

team can replicate. The result is a system that can monitor every transaction while focusing investigative resource where it matters.

Real-World Scenario: AML at Scale

A Tier-1 bank with \$2 trillion in annual payment flows deploys an ML transaction monitoring system integrated with network graph analytics and NLP-driven adverse media screening. False positive rates drop from 96% to 45%, releasing 2,000 analyst hours per week for genuine investigation. The ML model is documented in the model risk management inventory, validated by an independent model validation unit, and monitored quarterly for performance drift. The institution submits 40% more Suspicious Activity Reports (SARs) of materially higher quality. This is the governance-conditioned AI-GRC outcome the SEM validates.

6.2 AI in Credit Risk – From Scorecards to Dynamic Underwriting

The Underwriting Revolution

Traditional credit underwriting relies on structured data credit bureau scores, income verification, debt-to-income ratios applied through logistic regression scorecards that were often built on years of old data and calibrated for economic conditions that no longer exist. These scorecards exclude significant portions of the creditworthy population including borrowers who have deteriorated since their last bureau update. For lenders, this means adverse selection, suboptimal pricing, and loss rates that surprise on the downside when economic conditions shift.

ML-based credit risk models expand the data universe by incorporating transaction behavior, payment patterns, digital footprint data, macroeconomic indicators, and alternative data sources alongside traditional bureau data. They update continuously rather than waiting for annual model redevelopment cycles. And they detect non-linear relationships and interaction effects that logistic regression systematically misses. The AI-Wasleh case documented by Shaban and Omoush (2025) illustrates the practical implementation: ML algorithms analyzing transaction history, payment behavior, and socio-economic indicators generate real-time credit assessments that reduce human bias in loan approvals and produce more accurate risk stratification across the customer portfolio.

For financial institutions, the GRC implications extend beyond credit quality: Fair lending regulations (Equal Credit Opportunity Act in the US, Consumer Credit Directive in Europe) require that credit decisions be explainable and non-discriminatory, precisely the explainability challenge that Artene et al. (2024) identify as a structural limitation of neural network models.

6.3 AI in Regulatory Reporting and Compliance Automation

From Manual Consolidation to Intelligent Reporting

Regulatory reporting represents one of the highest-volume, most error-prone, and least value-generating activities in financial institution operations. A major bank's Regulatory Reporting team may produce hundreds of reports annually for the Federal Reserve, FINRA, the SEC, the ECB, national competent authorities, and tax jurisdictions; each with its own data definitions, calculation methodologies, submission formats, and deadlines. The data underlying these reports typically lives in dozens of source systems, requires manual extraction, transformation, and reconciliation, and is subject to amendment and restatement processes when errors are discovered. The Thomson Reuters Cost of Compliance 2023 report documents that financial services firms reviewed over 56,000 regulatory documents globally in 2022 with an average of more than 150 per day.

NLP-driven regulatory change management tools address this volume problem directly, parsing regulatory publications, extracting requirement changes, mapping them to internal policies and controls, and generating gap analyses. ML-based data quality scoring identifies reconciliation breaks and data lineage issues before they propagate into submitted reports. The JP Morgan COiN platform, demonstrates the scale of what NLP can achieve in legal and compliance document review; 360,000 hours of annual processing reduced through AI. Applied to regulatory reporting, the same capability class enables compliance teams to redirect analyst capacity from data wrangling to regulatory interpretation and decision-making.

Artene et al. (2024) observe that embedded AI systems with sophisticated algorithms demonstrate the ability to continuously learn from data streams, adapt to evolving patterns, and make real-time decisions by shifting control from deterministic models to adaptive probabilistic

frameworks. In the regulatory reporting context, this means systems that can detect when a regulatory requirement has changed and adjust extraction logic accordingly, rather than requiring a manual change management cycle that introduces latency and error risk between regulatory change and reporting response.

6.4 AI in Corporate Governance and Board-Level Transparency

Beyond Compliance - AI as a Governance Enabler

The governance application of AI extends beyond the compliance function into the boardroom. Shaban and Omoush (2025) provide the most granular empirical evidence on this dimension, with their survey of 564 financial sector professionals demonstrating that AI adoption significantly improves board-level decision quality ($R^2 = 0.582$), stakeholder engagement ($R^2 = 0.681$), and executive compensation committee effectiveness ($R^2 = 0.684$). These are not marginal improvements they represent strong structural effects that hold across firm size, sector, and AI adoption stage.

The practical mechanism is straightforward. Boards make better decisions when they have better information, and for most of the history of corporate governance that better information has been a management privilege. AI changes that asymmetry. Analytics platforms give board members real-time visibility into financial performance, risk concentrations, regulatory standing, and ESG metrics on the same timeline that management has it. AI-assisted internal audit surfaces patterns and anomalies that sampling-based human review will systematically miss – not because auditors are incompetent, but because sampling is structurally limited in what it can find. Predictive analytics on credit portfolios and market exposures give boards forward-looking information rather than the retrospective summaries they have traditionally received. These are governance improvements that do not require new legislation or new board structures; they require the right technology deployed with the right data.

Decision Harbor AI model provides an instructive architectural illustration of how AI can support board-level strategic decision-making. Their game-theoretic payoff matrix demonstrates that the optimal outcome of High/High for both management and shareholders require both parties to commit to AI-driven financial

reporting. When management adopts and shareholders resist (or vice versa), payoffs are asymmetric. This models a real governance dynamic: AI transparency initiatives require stakeholder buy-in across the principal-agent relationship to generate their full value.

7. Conclusion and future recommendations

7.1 Summary of Empirical Findings

Table 3: Hypothesis Results Summary

H	Hypothesis	Key Metric / Evidence	Paper	Verdict
H1	AI improves regulatory compliance effectiveness, mediated by governance maturity	$B (AI \rightarrow Gov) = 0.76$; $\beta (Gov \rightarrow Outcomes) = 0.73$; both $p < 0.001$	Choowan and Daovisan	Supported
H2	AI risk analytics outperform rule-based systems, conditional on data governance quality	Fraud detection MSVR 0.81; Credit scoring MSVR 0.78; Quantile invariance $\rho = 0.91$	Choowan and Daovisan	Supported
H3	AI-enabled transparency strengthens board decision quality, stakeholder engagement and audit effectiveness	R^2 : decision-making 0.582; stakeholder engagement 0.681; executive compensation 0.684; all $p < 0.001$	Shaban and Omoush	Supported
H4	AI-GRC returns are higher for advanced governance maturity institutions	Compliance automation $\beta = 0.63$ (Q0.25) vs 0.86 (Q0.75); fraud detection $\beta = 0.57$ vs 0.81	Choowan and Daovisan	Supported

Taken together, the evidence reviewed in this paper points to four findings that are worth stating directly, because they challenge some of the assumptions that still govern how financial institutions approach AI-GRC investment.

- 1. Governance is not the bureaucratic overhead of AI deployment – it is the determinant of AI performance:** The SEM evidence is unambiguous: governance maturity of **AI-to-outcome** relationship has a coefficient of 0.73. Institutions that treat data governance, model risk management, and compliance frameworks as costs to minimize are systematically under-extracting ROI from their AI investments. The quantile analysis makes this concrete: compliance automation performs 37% better at advanced governance maturity than at nascent maturity, and fraud detection 42% better. Those differentials do not close by buying a better AML system. They close by building the data infrastructure and model governance that the current system cannot utilize.
- 2. Machine learning and NLP, not expert systems or rules engines, are the AI capabilities that drive GRC value:** ML carries an MSVR weight of 0.87, which is the highest of any AI technique across the governance performance spectrum. NLP follows at 0.74. Expert systems and fuzzy logic score 0.59 and 0.41

respectively. For institutions still investing in rule-based compliance strategy as their primary AI strategy, the evidence suggests a misallocation of resource relative to what the market has validated.

- 3. Fraud detection and AML are where AI delivers the most immediate and quantifiable GRC value:** Fraud detection carries the highest outcome weight (MSVR 0.81) and the steepest maturity gradient from Q0.25 to Q0.75. Credit scoring follows at MSVR 0.78. These domains are where the volume-to-human capacity mismatch is most severe, the data characteristics most favorable to ML, and the regulatory tolerance for false positives lowest. They are also where the business case for AI investment is most straightforwardly expressed in dollar terms.
- 4. AI improves governance at the corporate level, not just the operational level:** Shaban and Omoush (2025) document AI's effect on board decision quality, stakeholder engagement, and executive accountability; governance dimensions that sit above the operational compliance layer. AI is not just a tool for processing transactions faster; it is a mechanism for improving the information environment in which principals (shareholders, boards) oversee agents (management, risk-takers). The Agency Theory prediction that better information reduces

governance costs is empirically confirmed in financial market institutions that have integrated AI into their transparency and reporting processes.

7.2 Recommendations

Study carries out some of the recommendations for smooth AI-GRC adoption in financial markets




Fig 3: Strategic framework for AI-driven GRC implementation

<https://aimjournals.com/index.php/ijmcsit>78

- **Adopt a governance-first AI implementation approach:** Start with a governance audit, not a vendor shortlist. Before committing capital to AI tools, map the current data governance state against the dimension's SEM evidence identifies as determinative: data quality, catalog coverage, lineage tracking, and model risk management maturity. Identify the gaps that will constrain AI GRC performance before performing vendors platform feature comparison.
- **AI-Governance investment approach:** Sequence your AI-GRC investment to match governance maturity. The evidence shows that governance investment pays a higher margin return than incremental AI capability investment at early maturity stages. Compliance automation and data quality management deliver meaningful returns even at nascent implementation ($\beta = 0.63, 0.58$ at Q0.25).
- **Establish an AI GRC Center of Excellence:** Create a cross-functional team encompassing model risk management, compliance, data governance, and technology to govern the full lifecycle of GRC AI models from development through deployment, monitoring, and decommissioning. This institutional structure is the organizational manifestation of the governance maturity that the data demonstrates is decisive.
- **Prioritize explainability of AI-GRC:** Build explainability (XAI) into AI-GRC architecture from the outset, not as an afterthought. For any AI model used in credit decisions, AML adjudication, or regulatory reporting, designate the explainability methodology as part of the model specification, not the implementation review.
- **Implement perpetual KYC using AI:** Establish perpetual KYC using AI-driven continuous monitoring rather than periodic refresh. Customers whose risk profiles are monitored continuously through transaction behavior analytics, adverse media screening, and sanctions list updates present materially lower residual risk than customers reviewed annually on a calendar basis.
- **Reframe AI-GRC strategy:** Elevate AI-GRC from a technology program to a strategic governance investment, not as a cost center. The empirical evidences on board decision quality, and on stakeholder engagement demonstrates that AI's governance effects operate at the board level, not just the operational level.
- **Mandate AI ethics and Governance frameworks:** Establish corporate AI ethics committees to institutionalize algorithmic accountability and enforce transparency. Designate AI governance boards within organizational structures to ensure that AI-driven financial decisions align with ethical, legal, and regulatory standards reflecting the XAI governance recommendations of Shaban and Omoush (2025).
- **Invest in AI literacy:** Invest in AI literacy across the first and second lines of defense. The AI-Wasleh case documents how staff resistance to automated decision-making was a major implementation barrier, requiring AI literacy programs and role restructuring before AI-GRC could be effectively operated. AI literacy programs for compliance officers, risk managers, and internal auditors are a prerequisite for effective human-AI collaboration in GRC functions.
- **Update Model Risk Management guidance:** Develop AI-specific model risk management guidance that addresses the unique characteristics of ML models (non-linearity, feature drift, distributional shift) beyond the current framework designed for traditional statistical models.
- **Harmonize cross-jurisdictional AI governance standards:** The tension between GDPR's data minimization requirements and ML's data intensity demands is unresolved and creates compliance complexity for multinational institutions operating across the EU and other jurisdictions. International regulatory coordination through the FSB and Basel Committee is needed to establish coherent global AI-GRC standards.

7.3 Summary and Core Insights

Table 4: Core Insights and Recommended Actions

#	Insight	What Institutions Should Do
1	Governance maturity is the decisive variable in AI-GRC performance not algorithm sophistication	Invest in data governance infrastructure, model risk frameworks, and compliance automation before expanding AI vendor relationships
2	ML and NLP deliver measurably superior GRC returns versus rule-based systems across all major use cases	Migrate AML, credit risk, and regulatory reporting to ML/NLP architectures; treat rule engines as legacy infrastructure to phase out
3	AML and fraud detection generate the highest immediate AI-GRC value and the steepest maturity gradient	Prioritize financial crime surveillance as the primary AI-GRC entry point; build governance around it before expanding to adjacent domains
4	Explainability is the critical unresolved challenge for AI in highly regulated financial decision-making	Designate explainability (XAI) methodology as a first-class model specification requirement; engage XAI specialists during model design, not post-deployment
5	AI improves corporate governance quality at the board and executive level, not just operational compliance	Extend AI transparency tools to stakeholder dashboards, board reporting, and ESG disclosure. Treat governance AI as a strategic, not just a compliance investment
6	Cultural and organizational readiness are as important as technical capability	Fund AI literacy programs for compliance, risk and audit staff; build governance committees with genuine authority; invest in change management alongside technology

8. References

Reference Papers

1. Choowan, P., & Daovisan, H. (2026). *Artificial Intelligence in Data Governance for Financial Decision-Making: A Systematic Review.* *Big Data and Cognitive Computing*, 10(1), 8
2. Artene, A. E., Domil, A. E., & Ivascu, L. (2024). *Unlocking Business Value: Integrating AI-Driven Decision-Making in Financial Reporting Systems.* *Electronics*, 13(15), 3069.
3. Shaban, O. S., & Omoush, A. (2025). *AI-Driven Financial Transparency and Corporate Governance: Enhancing Accounting Practices with Evidence from Jordan.* *Sustainability*, 17(9), 3818.
4. Avelar, E. A., & Jordao, R. V. D. (2025). *The role of artificial intelligence in the decision-making process: A study on the financial analysis and movement forecasting of the world's largest stock exchanges.* *Management Decision* 63 (10): 3533–3556.
5. DiMaggio, P. J., & Powell, W. W. (1983). *The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields.* *American Sociological Review*, 48, 147–160.
6. Easley, D., & O'Hara, M. (2005). *Information and the cost of capital.* *Journal of Finance*, 59, 1553–1583.
7. Fedyk, A., Hodson, J., Khimich, N., & Fedyk, T. (2022). *Is artificial intelligence improving the audit process? Review of Accounting Studies*, 27, 938–985.
8. Gyau, E. B., Appiah, M., Gyamfi, B. A., Achie, T., & Naeem, M. A. (2024). *Transforming banking: Examining the role of AI technology innovation in boosting banks' financial performance.* *International Review of Financial Analysis*, 96, 103700.
9. Hair, J. F., Astrachan, C. B., Moisescu, O. I., Radomir, L., Sarstedt, M., Vaithilingam, S., & Ringle, C. M. (2021). *Executing and interpreting applications of PLS-SEM: Updates for family business researchers.* *Journal of Family Business Strategy*, 12, 100392.
10. Jensen, M. C., & Meckling, W. H. (1976). *Theory of the firm: Managerial behavior, agency costs, and ownership structure.* *Journal of Financial Economics*, 3, 305–360.
11. Manta, O., Vasile, V., & Rusu, E. (2025). *Banking Transformation Through FinTech and the Integration of Artificial Intelligence in Payments.* *FinTech*, 4(2), 13.
12. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). *A survey on bias and fairness in machine learning.* *ACM Computing Surveys*, 54, 1–35.

13. Omoteso, K., & Mobolaji, H. (2020). *AI and Corporate Governance: Enhancing Transparency in Financial Reporting*. *Journal of Business Ethics*, 165(4), 789–805.
14. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., & Moher, D. (2021). *The PRISMA 2020 statement: An updated guideline for reporting systematic reviews*. *BMJ*, 372, n71.
15. Sai, S., Arunakar, K., Chamola, V., Hussain, A., Bisht, P., & Kumar, S. (2025). *Generative AI for finance: Applications, case studies and challenges*. *Expert Systems*, 42(3), e70018.
16. Scott, W. R. (2014). *Institutions and Organizations: Ideas, Interests, and Identities* (4th ed.). SAGE Publications.
17. Weber, P., Carl, K. V., & Hinz, O. (2024). *Applications of explainable artificial intelligence in finance*. *Management Review Quarterly*, 74, 867–907.
3. European Banking Authority. (2021). *Guidelines on Internal Governance* (EBA/GL/2021/05).
4. Federal Reserve Board. (2011). *Supervisory Guidance on Model Risk Management* (SR Letter 11-7). Board of Governors of the Federal Reserve System.
5. Financial Action Task Force. (2023). *FATF Guidance on Risk-Based Approach to Virtual Assets*. FATF-GAFI.
6. Financial Stability Board. (2022). *Artificial Intelligence and Machine Learning in Financial Services*. FSB.
7. FinCEN. (2022). *FinCEN Issues Statement on the Use of Artificial Intelligence in AML Compliance Programs*. US Department of the Treasury.
8. McKinsey & Company. (2023). *AI-Enabled Compliance: From Regulatory Burden to Strategic Advantage*. McKinsey Financial Services Practice.

Regulatory and Industry References

1. Basel Committee on Banking Supervision. (2021). *Principles for Operational Resilience*. Bank for International Settlements.
2. Boston Consulting Group. (2023). *Global Financial Services Regulatory Fines Report 2023*. BCG Financial Institutions Practice.
9. PricewaterhouseCoopers. (2021). *The Future of AI in Corporate Governance: Implications for Risk and Compliance*. PwC Insights.
10. Thomson Reuters. (2023). *Cost of Compliance 2023: Shaping the Future of Financial Services Compliance*. Thomson Reuters Regulatory Intelligence.