INTERNATIONAL JOURNAL OF ARTS, COMMERCE & EDUCATION
A Multidisciplinary Peer-Reviewed, Refereed, Open Access Journal,
https://avjournals.com/ e-ISSN: 2347-7687, Vol 13, Issue 2, July-December 2025

INTEGRATING AI-ENABLED PREDICTIVE ANALYTICS, FORENSIC INVESTIGATION, AND ENTERPRISE RISK MANAGEMENT FOR EFFECTIVE CORPORATE SECURITY RISK GOVERNANCE

[1]Amit Dattatray Patil, [2]Dr. Gulab Singh
[1]Research Scholar, [2]Research Supervisor
[1-2]Department of Management, Arunodaya University, Itanagar, Arunachal Pradesh

Abstract
Corporate security risk is usually handled by three functions that seldom work as one: security analytics teams that detect anomalies, forensic investigators who establish what happened, and enterprise risk management (ERM) functions that report exposure to the board. This paper argues that effective governance depends less on the strength of any one of these functions than on the links between them, and analyses those links using published industry evidence and simple quantitative models. Breach data show that the global average cost of a data breach was USD 4.44 million in 2025 and that organizations making extensive use of security AI and automation reported costs USD 1.9 million lower and breach lifecycles 80 days shorter. Fraud data from 1,921 investigated cases show a median loss of USD 145,000 and a typical duration of twelve months before detection; frauds detected by automated monitoring had a median loss of USD 83,000 and lasted six months, against USD 227,000 and eighteen months for those found by external audit. Re-analysis shows that loss accumulates at a broadly similar monthly rate whatever the detection method, so that detection time, not detection method as such, drives loss. A base-rate analysis shows the main limit of predictive analytics: with one malicious event in a thousand, a detector that catches 95 per cent of incidents and wrongly flags 1 per cent of normal events produces alerts of which only 8.7 per cent are real. Analytics therefore creates value only when it is coupled to investigative capacity that can triage alerts and to a feedback loop that returns confirmed cases as training data. The paper proposes an integrated governance framework built on COSO ERM, ISO 31000 and the Govern function of NIST CSF 2.0, specifies the information flows between the three functions, discusses the evidentiary and model-risk requirements that AI introduces, and sets out metrics for boards. It concludes that integration is the main source of value and that the weakest link is usually the passage from alert to evidence.
Keywords: Corporate security; risk governance; enterprise risk management; predictive analytics; artificial intelligence; digital forensics; fraud detection; model risk

1. Introduction

Security risk has moved from the server room to the boardroom. A single breach now costs a large organization millions in direct expenses and more in lost trust, regulatory penalties and disruption. Fraud by insiders, the oldest form of corporate security risk, continues to drain a substantial share of revenue. Regulators have responded by making boards answerable for these risks. In the United States, listed companies must disclose a material cybersecurity incident within four business days of determining that it is material. In India, the Computer Emergency Response Team requires specified incidents to be reported within six hours of being noticed. Governance of security risk has become a legal duty as well as a management task.

Organizations have answered with three sets of capabilities. Predictive analytics, increasingly built on machine learning, scores events, users and transactions for the likelihood of malicious activity. Forensic investigation establishes, after an alert or an incident, what happened, who was responsible and what evidence supports that conclusion. Enterprise risk management places security risk alongside financial, operational and strategic risks, sets an appetite for it and reports it to the board. Each of these has its own professional community, standards and tools.

The difficulty is that they are often run in isolation. Analytics teams generate alerts that no one has time to investigate. Investigators work on cases whose lessons never reach the models or the risk register. Risk managers report exposure using qualitative heat maps that do not reflect either the alerts or the findings. In such an organization each function may be competent and the whole may still fail, because the failures occur in the hand-offs between them.

This paper examines that problem analytically. Its objectives are:

- to set out the contribution and limits of each of the three functions in corporate security risk governance;
- to quantify, from published evidence, the value of early and automated detection;
- to analyse the base-rate limit of predictive analytics and its consequence for investigative capacity;
- to propose an integrated framework and specify the information flows between the functions;

INTERNATIONAL JOURNAL OF ARTS, COMMERCE & EDUCATION
A Multidisciplinary Peer-Reviewed, Refereed, Open Access Journal,
https://avjournals.com/ e-ISSN: 2347-7687, Vol 13, Issue 2, July-December 2025

- to identify the evidentiary, legal and model-risk requirements introduced by the use of AI; and
- to propose metrics by which a board can judge whether the integration works.

2. Background

2.1 Predictive analytics in security

Security analytics has moved from fixed rules, such as an alert on more than five failed logins, to statistical and machine-learning models that learn what normal behaviour looks like and flag departures from it. User and entity behaviour analytics builds profiles of individual users and devices; fraud analytics scores transactions; threat intelligence platforms forecast which vulnerabilities are likely to be exploited. Two developments have accelerated this. Volumes of log and transaction data have grown far beyond what people can read, and generative AI has begun to assist analysts in summarising alerts and writing queries.

The evidence on value is encouraging but needs careful reading. In the 2025 study of 600 breached organizations conducted for IBM, the global average cost of a breach was USD 4.44 million, the first decline in five years, and the average time to identify and contain a breach was 241 days. Organizations that used security AI and automation extensively reported average costs USD 1.9 million lower and breach lifecycles 80 days shorter than those that did not (IBM, 2025). The same study found that 13 per cent of organizations reported breaches of AI models or applications, that 97 per cent of those lacked proper AI access controls, and that 63 per cent of breached organizations had no AI governance policy or were still developing one. AI is thus both a tool for managing security risk and a new source of it.

2.2 Forensic investigation

Forensic investigation answers four questions: what happened, how, who was responsible and what can be proved. In digital incidents it involves identifying, collecting, preserving and analysing electronic evidence in a way that keeps its integrity demonstrable, which international guidance sets out in terms of identification, collection, acquisition and preservation (ISO/IEC 27037:2012). NIST guidance describes how forensic techniques should be built into incident response rather than added after it (Kent et al., 2006). In fraud, investigation combines document and transaction analysis with interviews and often leads to disciplinary, civil or criminal proceedings.

Investigations are the main source of reliable knowledge about how attacks and frauds succeed. The 2025 Data Breach Investigations Report analysed 22,052 incidents, of which 12,195 were confirmed breaches. Breaches involving a third party doubled from 15 to 30 per cent, exploitation of vulnerabilities accounted for 20 per cent of initial access, ransomware was present in 44 per cent of breaches, and the human element was involved in about 60 per cent (Verizon, 2025). Figures of this kind exist only because cases were investigated to a conclusion.

2.3 Enterprise risk management

Enterprise risk management treats risk as a matter for the whole organization and its strategy. The COSO framework of 2017 organizes ERM into five components and twenty principles, from governance and culture to review and revision (COSO, 2017), and ISO 31000 provides a general process of risk identification, analysis, evaluation, treatment, monitoring and communication (ISO, 2018). The Three Lines Model assigns responsibility to management, to risk and compliance functions, and to internal audit (IIA, 2020). In 2024 version 2.0 of the NIST Cybersecurity Framework added a sixth function, Govern, to the existing Identify, Protect, Detect, Respond and Recover, in recognition that cybersecurity is a major source of enterprise risk for senior leaders to consider alongside financial and reputational risk (NIST, 2024). Table 1 summarises what each framework contributes to the integration considered here.

Table 1. Frameworks relevant to integrated security risk governance

Framework	Scope	Contribution to integration	What it leaves open
COSO ERM (2017)	Enterprise risk and strategy	Board oversight, risk appetite, link to performance	No technical guidance on detection or evidence
ISO 31000 (2018)	Generic risk process	Common process and vocabulary for all risk types	No security-specific controls
NIST CSF 2.0 (2024)	Cybersecurity outcomes	Govern function ties cyber risk to enterprise risk	Fraud and physical security outside scope
NIST AI RMF 1.0 (2023)	Risks of AI systems	Govern, Map, Measure, Manage functions for models	Not specific to security use of AI
ISO/IEC 27037 (2012)	Digital evidence handling	Integrity and chain of custody of evidence	No link to analytics or risk reporting
Three Lines Model (2020)	Assurance roles	Separates ownership, oversight and audit	Does not specify information flows

INTERNATIONAL JOURNAL OF ARTS, COMMERCE & EDUCATION
A Multidisciplinary Peer-Reviewed, Refereed, Open Access Journal,
https://avjournals.com/ e-ISSN: 2347-7687, Vol 13, Issue 2, July-December 2025

2.4 The gap

Each framework addresses part of the problem. None specifies how the outputs of predictive models should feed investigations, how investigative findings should change models and risk assessments, or how either should be expressed in the terms of enterprise risk. This gap between functions is where the analysis of this paper is directed.

3. Method

The paper is an analytical study based on secondary evidence. No primary data were collected. Three kinds of input were used.

- **Industry studies with stated methods.** The IBM Cost of a Data Breach Report 2025, conducted by the Ponemon Institute on 600 organizations; the Verizon Data Breach Investigations Report 2025, covering 22,052 incidents; and the Association of Certified Fraud Examiners (ACFE) Occupational Fraud 2024 report, based on 1,921 cases investigated by certified fraud examiners in 138 countries and territories.
- **Frameworks and standards.** COSO ERM, ISO 31000, NIST CSF 2.0, the NIST AI Risk Management Framework, ISO/IEC 27037 and supervisory guidance on model risk.
- **Quantitative models.** Simple models were applied to the published figures: loss accrual per month of undetected fraud, the precision of a detector under a given base rate, and annualised loss expectancy.

The precision of a detector, that is the share of its alerts that are real incidents, was calculated from its sensitivity s, its false-positive rate f and the prevalence π of malicious events:

$$\text{Precision} = (s \times \pi) / (s \times \pi + f \times (1 - \pi))$$

Annualised loss expectancy (ALE) was calculated as the single loss expectancy (SLE) multiplied by the annual rate of occurrence (ARO), and the return on security investment (ROSI) as the reduction in ALE less the annual cost of the control, divided by that cost:

$$\text{ALE} = \text{SLE} \times \text{ARO} \qquad \text{ROSI} = (\text{ALE}_{\text{before}} - \text{ALE}_{\text{after}} - \text{Cost}) / \text{Cost}$$

The industry figures are associations reported by the publishers, not results of controlled experiments, and the analysis treats them as such.

4. Results and Discussion

4.1 The value of early detection

The ACFE data allow the effect of detection on loss to be examined directly. Table 2 gives the median loss and median duration of frauds by the way in which they were detected, together with the loss per month of duration calculated for this paper.

Table 2. Median loss and duration of occupational fraud by method of detection

Detection method	Median loss (USD)	Median duration (months)	Loss per month of duration (USD)
Automated transaction or data monitoring	83,000	6	13,800
Tip	155,000	12	12,900
External audit	227,000	18	12,600
By accident	110,000	18	6,100

Frauds found by automated monitoring were detected in half the time of those found by tips and a third of the time of those found by external audit, and their median loss was correspondingly lower. The last column shows why. For the three systematic methods, the loss per month of duration lies in a narrow band between USD 12,600 and 13,800. Losses accumulate at a broadly similar rate whatever eventually reveals the fraud, so the loss at discovery is governed mainly by how long the fraud ran. Frauds found by accident accumulate losses at half the rate, which suggests that they are smaller schemes that happen to surface rather than large ones caught early. The value of detection technology lies chiefly in shortening the period over which loss accumulates.

The same pattern appears for preventive controls. Organizations with proactive data monitoring and analysis had a median fraud loss of USD 100,000 against USD 200,000 without it, and a median duration of nine months against eighteen. Surprise audits, financial statement audits and hotlines were also associated with reductions of at least 50 per cent in both loss and duration, surprise audits showing the largest reduction in loss at 63 per cent (ACFE, 2024).

Two further findings point to the role of governance. A lack of internal controls contributed to 32 per cent of cases and the override of existing controls to 19 per cent, and in 84 per cent of cases the perpetrator displayed at least one behavioural red flag. Most frauds therefore leave signs that a well-informed analytics function could use, and most succeed through weaknesses that a risk function should already know about.

INTERNATIONAL JOURNAL OF ARTS, COMMERCE & EDUCATION
A Multidisciplinary Peer-Reviewed, Refereed, Open Access Journal,
https://avjournals.com/ e-ISSN: 2347-7687, Vol 13, Issue 2, July-December 2025

Tips remain the single largest source of detection, at 43 per cent of cases. Analytics complements this channel rather than replacing it, and an integrated system should treat tips as one more signal to be scored, triaged and investigated alongside model alerts.

4.2 The base-rate limit of predictive analytics

The second result concerns a limit that is mathematical rather than technical. Malicious events are rare among the millions of logins, transactions and file accesses that a large organization records each day. Table 3 shows what this means for a detector that identifies 95 per cent of malicious events, at two levels of prevalence and three false-positive rates, for one million events.

Table 3. Precision and alert volume of a detector with 95 per cent sensitivity (per million events)

Prevalence of malicious events	False-positive rate (%)	Real incidents detected	False alerts	Alerts that are real (%)
1 in 1,000	1	950	9,990	8.7
1 in 1,000	0.1	950	999	48.7
1 in 1,000	0.01	950	100	90.5
1 in 10,000	1	95	9,999	0.9
1 in 10,000	0.1	95	1,000	8.7
1 in 10,000	0.01	95	100	48.7

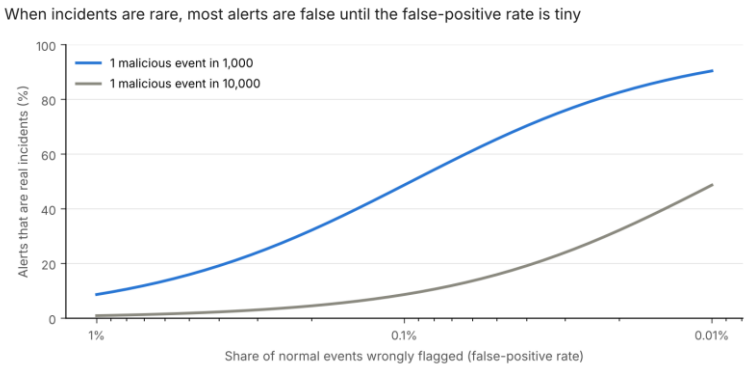

Figure 1. Share of alerts that are real incidents against the false-positive rate, at two levels of prevalence.

A false-positive rate of 1 per cent sounds small and would be considered good performance in many classification problems. At a prevalence of one in a thousand it means that more than nine alerts in ten are false; at one in ten thousand, ninety-nine in a hundred. The detector would generate about 10,000 alerts per million events, and an investigation team that can examine a few hundred cases a day would be overwhelmed. Only when the false-positive rate falls to 0.01 per cent does the majority of alerts become real, and that level is difficult to reach and to keep.

Three consequences follow for governance. First, a model must be judged by its precision at the organization's actual prevalence and by the alert volume it produces, not by its accuracy, which is dominated by the mass of normal events. Second, the capacity of the investigation function is a design parameter of the analytics function: the alert threshold should be set so that the volume of alerts matches what can be investigated, with the highest-scoring alerts first. Third, the investigation function is what turns a stream of probabilistic scores into knowledge. Every alert that is investigated and closed, as real or false, is a labelled example that can improve the model. Without that feedback the model cannot learn which of its alerts matter, and its false-positive rate will not fall.

The figures of the IBM study should be read with the same care. The lower cost and shorter lifecycle reported by organizations that use AI extensively, USD 1.9 million or 43 per cent of the average cost and 80 days or about a third of the average lifecycle, are differences between groups of organizations. Those that invest heavily in AI are likely also to have mature processes, skilled staff and integrated response, so not all of the difference can be attributed to the technology. The result is consistent with the argument of this paper: AI is associated with lower losses where it is part of a functioning detection-to-response chain.

4.3 An integrated framework

The analysis points to a framework in which the three functions are linked by explicit flows of information under a single governance layer. Figure 2 shows the framework.

INTERNATIONAL JOURNAL OF ARTS, COMMERCE & EDUCATION
A Multidisciplinary Peer-Reviewed, Refereed, Open Access Journal,
https://avjournals.com/ e-ISSN: 2347-7687, Vol 13, Issue 2, July-December 2025

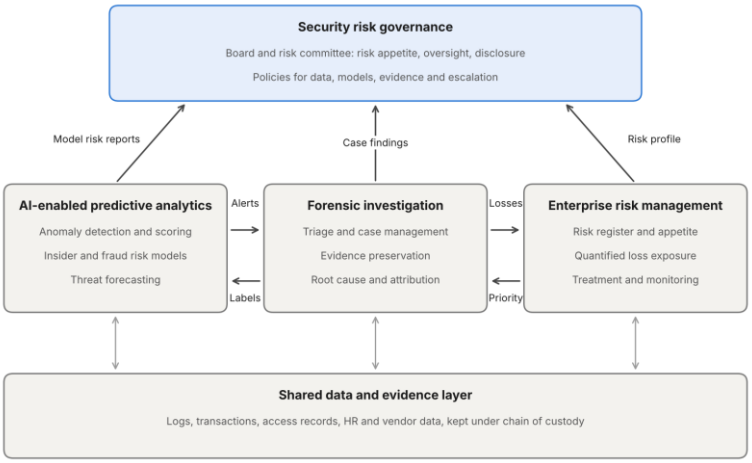

Figure 2. Integrated framework for corporate security risk governance.

Four flows connect the functions.

- **Alerts, from analytics to investigation.** Each alert carries a risk score, the evidence that produced it and an explanation of the main contributing factors. Alerts are ranked so that the investigation team works on the highest scores within its capacity.
- **Labels, from investigation to analytics.** Each closed case is returned with its outcome (confirmed, false or inconclusive) and the indicators that proved decisive. These labels are used to retrain and recalibrate the models.
- **Losses, from investigation to risk management.** Confirmed cases are recorded with their actual loss, root cause and the controls that failed. This builds the internal loss database on which quantified risk assessment depends.
- **Priorities, from risk management to investigation and analytics.** The risk register and risk appetite determine which scenarios the models should target and how investigative effort should be allocated between them.

Beneath the three functions lies a shared data and evidence layer. Logs, transactions, access records and human-resource and vendor data are collected once, retained according to policy and preserved with a chain of custody from the moment of collection. This is the principle of forensic readiness: an organization that collects data for analytics in a forensically sound way does not need to reconstruct evidence after an incident, and an investigation can proceed directly from an alert. Above the three functions, the governance layer sets risk appetite, approves policies for data, models and evidence, receives reports and is accountable for disclosure.

4.4 Quantifying exposure for the board

The framework allows security risk to be reported to the board in the same financial terms as other enterprise risks. A simple research uses the breach figures above. Suppose an organization judges that it faces a breach of average severity about once in five years, an annual rate of occurrence of 0.2. With a single loss expectancy of USD 4.44 million the annualised loss expectancy is about USD 0.89 million. If extensive use of AI and automation reduced the cost of a breach by the USD 1.9 million reported in the IBM study, the expectancy would fall to about USD 0.51 million, a reduction of USD 0.38 million a year. For a programme costing USD 0.25 million a year, the return on security investment would be about 52 per cent.

The result is sensitive to the assumed rate of occurrence. At one breach in ten years the annualised loss expectancy is about USD 0.44 million and at one in three to four years about USD 1.33 million, and the return on the same programme varies accordingly. A board should therefore receive the expectancy as a range with its main assumptions stated, not as a single figure.

The study is deliberately simple, and each number in it is uncertain. Its purpose is to show the form in which the integrated framework allows risk to be expressed. The rate of occurrence should come from the organization's own incident history and from sector data. The loss expectancy should come from the internal loss database built by the investigation function. The effect of the control should be estimated from the change in detection time observed after its introduction, as in Table 2, rather than taken from industry averages. A risk function that receives these inputs from the other two functions can replace qualitative heat maps by ranges of expected loss, which a board can compare directly with its risk appetite and with the cost of treatment.

4.5 Evidentiary and legal requirements

The use of AI creates requirements that do not arise with rule-based detection. An alert produced by a model may lead to dismissal, a civil claim or a criminal complaint. The evidence that supports it must therefore satisfy the standards of the forum in which it will be tested.

INTERNATIONAL JOURNAL OF ARTS, COMMERCE & EDUCATION
A Multidisciplinary Peer-Reviewed, Refereed, Open Access Journal,
https://avjournals.com/ e-ISSN: 2347-7687, Vol 13, Issue 2, July-December 2025

- **Integrity of electronic records.** Courts require proof that electronic evidence has not been altered. In India the Bharatiya Sakshya Adhiniyam, 2023 requires a certificate for the admission of electronic records, continuing the requirement previously found in Section 65B of the Indian Evidence Act. Data that were collected for analytics without hashing, time stamping and documented custody may be difficult to use.
- **Explainability of the model.** A model score is not in itself evidence of misconduct. The investigator must be able to show which observed facts led to the alert and must establish the misconduct from those facts by ordinary means. Models whose outputs cannot be explained in terms of observable behaviour are suitable for prioritising work but not for supporting conclusions.
- **Privacy and proportionality.** Monitoring of employees and customers is subject to data protection law, including India's Digital Personal Data Protection Act, 2023. The purpose, scope and retention of monitoring should be set by policy, approved under the governance layer and communicated to those affected.
- **Timeliness of reporting.** Reporting obligations, such as the six-hour requirement of CERT-In and the four-business-day disclosure rule for material incidents in the United States, run from the moment an incident is identified. An integrated framework in which analytics, investigation and risk assessment share data shortens the time needed to establish what has happened and whether it is material.

4.6 Model risk

A predictive model is itself a source of risk and must be governed as one. Supervisory guidance on model risk in banking requires independent validation, ongoing monitoring and clear ownership of models (Board of Governors of the Federal Reserve System, 2011), and the NIST AI Risk Management Framework sets out comparable functions for AI systems in general (NIST, 2023). Three risks are specific to security models.

- **Drift.** Normal behaviour changes with new systems, new working patterns and seasonal cycles, so a model's false-positive rate rises unless it is recalibrated. The label flow from investigation is the main defence.
- **Adversarial adaptation.** Attackers and fraudsters learn the boundaries of detection and stay inside them. A model that is never challenged by red-team exercises or new case findings will gradually detect less.
- **Bias and unfairness.** A model trained on past investigations reproduces any bias in whom the organization chose to investigate. Outcomes should be reviewed for disproportionate effects on particular groups of employees or customers.

The 2025 breach data add a fourth concern. AI systems used by the organization for other purposes are themselves targets, and in the study 97 per cent of organizations that reported AI-related breaches lacked proper access controls for AI. Governance of the security use of AI cannot be separated from governance of AI in general.

4.7 Metrics for the board

Integration has to be visible in the reports the board receives. Table 4 proposes a small set of metrics, one or two for each link of the framework.

Table 4. Board-level metrics for integrated security risk governance

Link in the framework	Metric	What it reveals
Analytics to investigation	Precision of alerts at actual prevalence; alerts per investigator per day	Whether analytics produces work that can be done
Investigation	Median time from alert to case conclusion; share of alerts closed	Whether alerts become knowledge
Investigation to analytics	Share of closed cases returned as labels; time since last recalibration	Whether models learn from cases
Investigation to risk management	Share of confirmed cases with loss and root cause recorded	Whether the loss database is complete
Risk management	Annualised loss expectancy against risk appetite, as a range	Whether exposure is within tolerance
Outcome	Median time to detect and contain; median loss per incident	Whether the whole system is improving

The metrics are deliberately few. Each measures a hand-off rather than the activity of a single function, because the analysis has shown that the hand-offs are where value is lost.

4.8 Barriers to integration

Several obstacles explain why integration is rare in practice.

INTERNATIONAL JOURNAL OF ARTS, COMMERCE & EDUCATION
A Multidisciplinary Peer-Reviewed, Refereed, Open Access Journal,
https://avjournals.com/ e-ISSN: 2347-7687, Vol 13, Issue 2, July-December 2025

- **Organizational separation.** Analytics often sits in the security operations centre, investigation in legal, internal audit or a fraud unit, and risk management in finance or a risk office, each reporting to a different executive.
- **Incompatible data.** The three functions use different case identifiers, taxonomies and systems, so that a case cannot easily be followed from alert to loss.
- **Confidentiality of investigations.** Investigation findings are often privileged or sensitive and are not shared with analytics teams, which deprives the models of labels.
- **Scarcity of skills.** Few staff combine data science, investigative and risk expertise, and the functions do not speak each other's language.
- **Qualitative risk culture.** Where risk is reported in heat maps, there is no place to put quantitative evidence from analytics and investigation, and so no demand for it.

Each barrier corresponds to an element of the framework: a single governance layer addresses separation, a shared data layer addresses data, controlled sharing of labels without case details addresses confidentiality, joint teams address skills, and quantified reporting addresses culture.

4.9 A staged path to integration

Few organizations can build the full framework at once. The analysis suggests a sequence of four stages, each of which creates the conditions for the next.

1. **Common records.** Agree a single case identifier and taxonomy for alerts, investigations and risk events, and collect key logs in a forensically sound way. Without this, no flow between the functions can be measured.
2. **Managed alert flow.** Set alert thresholds to match investigative capacity, rank alerts by score and record the outcome of every investigated alert.
3. **Closed learning loop.** Return investigation outcomes to the models as labels, recalibrate on a fixed schedule, and record confirmed losses and root causes in the risk register.
4. **Quantified governance.** Report security risk to the board as ranges of expected loss against appetite, using the internal loss database and the observed effect of controls on detection time.

The stages correspond to the metrics of Table 4, so that a board can see which stage an organization has reached and what it needs to do next.

5. Conclusion

This paper has analysed how AI-enabled predictive analytics, forensic investigation and enterprise risk management can be integrated for the governance of corporate security risk. The main conclusions are as follows.

1. **Detection time drives loss.** Frauds detected by automated monitoring ran for six months with a median loss of USD 83,000, against eighteen months and USD 227,000 for those found by external audit, and losses accumulated at a similar monthly rate across the systematic detection methods.
2. **Predictive analytics is limited by base rates.** At realistic prevalence most alerts are false unless the false-positive rate is extremely low, so analytics creates value only when coupled to investigative capacity and to feedback from investigated cases.
3. **Investigation is the link that turns alerts into knowledge.** It supplies labels to the models and loss data to the risk function, and it is the most common weak point.
4. **Enterprise risk management gives the board a common measure.** Fed by the other two functions, it can express security risk as expected loss and compare it with appetite and with the cost of controls.
5. **AI brings new obligations.** Evidence must be preserved in a forensically sound way from the start, models must be explainable, validated and monitored for drift and bias, and AI systems must themselves be protected.

The value of the three functions lies mainly in the links between them. An organization that strengthens any one function without connecting it to the others is likely to see less benefit than one that builds the flows of alerts, labels, losses and priorities described here under a single governance layer.

Limitations and future work. The study relies on industry reports whose samples are not random and whose findings are associations, not causal effects. The fraud and breach data come from different populations and years. The quantitative models are simple, and the return calculation. The framework has not been tested in an organization. Future research should evaluate the framework through case studies, measure the effect of integration on detection time and loss within organizations, and examine how the requirements of Indian evidence and data protection law apply to AI-generated alerts.

References

- Association of Certified Fraud Examiners. (2024). *Occupational fraud 2024: A report to the nations*. ACFE.
- Board of Governors of the Federal Reserve System. (2011). *Supervisory guidance on model risk management* (SR Letter 11-7).
- Committee of Sponsoring Organizations of the Treadway Commission. (2017). *Enterprise risk management: Integrating with strategy and performance*. COSO.
- IBM. (2025). *Cost of a data breach report 2025*. IBM Corporation and Ponemon Institute. https://www.ibm.com/reports/data-breach
- Institute of Internal Auditors. (2020). *The IIA's Three Lines Model: An update of the Three Lines of Defense*. IIA.
- International Organization for Standardization. (2012). *ISO/IEC 27037:2012 Information technology — Security techniques — Guidelines for identification, collection, acquisition and preservation of digital evidence*. ISO.
- International Organization for Standardization. (2018). *ISO 31000:2018 Risk management — Guidelines*. ISO.
- Kent, K., Chevalier, S., Grance, T., & Dang, H. (2006). *Guide to integrating forensic techniques into incident response* (NIST Special Publication 800-86). National Institute of Standards and Technology.
- National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). NIST.
- National Institute of Standards and Technology. (2024). *The NIST cybersecurity framework (CSF) 2.0* (NIST CSWP 29). NIST.
- Verizon. (2025). *2025 data breach investigations report*. Verizon Business.