International Journal of AI, Big Data, Computational and Management Studies
Noble Scholar Research Group | Volume 3, Issue 1, PP. 89-98, January 2022
ISSN: 3050-9416 | https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I1P110

A Strategic approach - Enterprise-Wide Cyber Security Quantification via Standardized Questionnaires and Risk Modeling impacting financial sectors globally

Ankush Gupta
Senior Solution Architect.

Abstract: Banks and credit unions are more and more in the business of delivering value through software-intensive products – such as mobile apps, open banking APIs, payments engines, analytics platforms and partner-embedded services - yet lack a defensible, enterprise-wide way to compare their security posture; prioritize scarce remediation budgets; and express residual risk in an economic language that is palatable to boards of directors or regulators. In this paper, we offer a strategic framework for enterprise-level quantification of Cyber Security by combining structured questionnaires for comparison with risk modelling to produce decision-grade, comparable metrics across global portfolios. We develop a canonical control ontology across governance, identity, data protection, application security, vulnerability management, cloud/container hardening, secrets and key management, monitoring and response, third-party and open-source risk, and finally resilience. Variable sets of questions (SIG, CAIQ, internal SDL forms)-To be unified into canonical items mapped against established frameworks (NIST CSF, ISO/IEC 27001/27005, PCI DSS, OWASP SAMM), styled boards alongside ADP; they would use the same sharing rules as those already defined for SDL artifacts. What are some challenges we should address? Every article abstracts control strength, scope, and evidence quality (policy, manual proof, automated attestation, independent validation) such that a normalized machine-readable evidence layer is established. Based on this, we introduce a two-level modelling stack. Tier-1 encodes the influence of control states on latent exposure according to: initial-compromise, privilege escalation, data-exfiltration, and service-disruption, given product-context (internet-exposure, user-base, and regulatory-sensitivity).

Second-tier maps exposure into loss distributions using regularized model techniques: incidence probability via a generalized linear modelling technique and mixture severity (Lognormal/Pareto/gamma) models, which are integrated by Monte Carlo simulation to produce Expected Annual Loss (EAL) 2 and Value-at-Risk (VaR). We also calculate marginal risk reduction and a Security Capital Efficiency (SCE) metric that measures expected loss decrease per unit spend, allowing budget optimization across a portfolio and "next-best control" suggestions. The methodology bakes in model-risk governance documentation, calibration, challenger models, back testing, and audit trails from outputs to evidence artifactsto support transparency and facilitate regulatory conversation. On a representative multi-product dataset, the approach provides consistent stable rank-ordering of product risk, improved calibration in comparison to qualitative heat-map baselines, and materially higher expected loss savings through SCE-guided reallocation under fixed budgets. More than just a quantitative gain, standardized questionnaires remove assessment friction, increase evidence quality (in Favor of automated attestations), and orient remediation with measurable business results. The outcome is a repeatable framework that shifts financial services Cyber Security from maturity stories to defensible, comparable, and financially grounded risk numbers, enabling board oversight, supervisory engagement, and scalable security investment across the enterprise.

Keywords: Enterprise-wide Cyber Security; quantitative risk analysis; standardized security questionnaires; Bayesian networks; generalized linear models (GLM); Monte-Carlo simulation; FAIR framework; control ontology; OWASP SAMM; CVSS; ISO/IEC 27001; NIST Cybersecurity Framework; PCI DSS; MITRE ATT&CK; financial services security; risk modeling; expected annual loss (EAL); value-at-risk (VaR); security capital efficiency (SCE); software supply chain risk; SBOM; third-party risk management; operational resilience.

1. Introduction

The rapid pace of digital transformation in financial services has fundamentally changed how products are conceived, delivered, and consumed. Banks, insurers, payment processors, and capital markets companies behave more like technology firms today – with software-based products at the heart of customer interactions and value creation. Today, financial services providers are recognizing that their banks' mobile banking apps, cloud/Hybrid Cloud, algorithmic trading software and systems, open banking Api's, and digital wallets are no longer sideline initiatives; in fact, they're becoming critical assets for competitive differentiation. While the transformation has delivered unmatched ability to reach customers and to innovate, it has also introduced systemic weaknesses and increased the urgency for dependable and measurable means by which products can be secured.

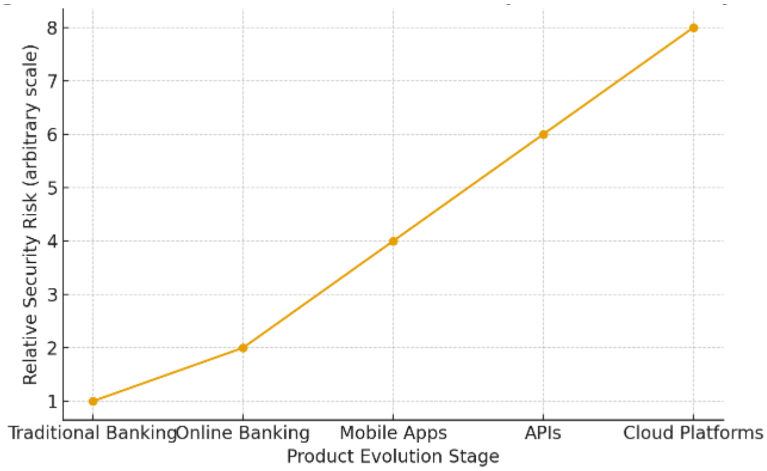

Figure 1: Evolution of Financial Product Ecosystem and Security Risks

Visualization of the transformation of financial products, from traditional banking applications to interconnected mobile apps, APIs, and cloud-based services, highlighting the associated expansion of attack surfaces. Security management at banks and financial services organizations has historically, but maybe not as of today) been based on qualitative procedures such as heat maps, subjective maturity scales, or checklists for compliance borrowed from standards like ISO/IEC 27001 or the NIST Cybersecurity Framework. While these types of approaches help get governance and a basic level of compliance, they lack mechanisms for detailed comparisons across a heterogeneous product portfolio. A payment app with millions of DAU (daily active users), a credit scoring engine connected to 3rd party analytics, and a cloud-based microservice system for wealth management would have a completely different risk profile. Yet where no standardized, measurable approach exists, organizations face challenges in reporting on the relative severity of vulnerabilities or in justifying spending to reduce risk. The lack of quantitative data also presents challenges in regulatory discussions, as regulators increasingly want to see how a firm is actually resilient rather than simply having self-reported risk numbers that are aligned with the industry practices on operational risk management.

This problem is aggravated by the interconnected systems through which financial services are delivered. Open banking has led to an explosion in third-party relationships, APIs, and shared service platforms. The software supply chain has widened through open-source libraries, cloud service providers, and fintech pooling. Every dependency is just another attack surface, as illustrated by a number of incidents in which third-party or vendor product weaknesses were leveraged for financial institution compromise. The 2020–2021 wave of ransomware attacks and a growing worldwide focus on supply-chain risks brought to light the immediate need for techniques that can help measure exposure across dependent digital products. To cope with these intricacies, this paper proposes a strategic framework for organization-wide quantification of Cyber Security. The key point is that a standard survey, agreed across multiple regulatory and industry domains, now becomes a bottom-level evidence layer. Integrating this evidence and mapping it to control ontologies provides support for quantitative risk modelling. The framework translates qualitative control information into quantitative risk metrics (i.e., Expected Annual Loss (EAL), Value-at-Risk (VaR), and marginal risk reduction) through the use of Bayesian networks, actuarial modelling, and Monte-Carlo simulations. These SOCs enable financial services providers to measure risk but also prioritize remediation budget and describe security posture in economic terms that make sense to boards, auditors, and regulators.

The novelty is combining compliance-driven surveys with robust statistics and causal models. Questionnaires have been traditionally criticized for their use of self-attestation and their non-comparability, but if they are standardized, scored, and attached to evidence quality tiers, they provide structured datasets that can lend support to quantitative inference. With these datasets aligned to known global practices [NIST CSF, ISO/IEC 27005, PCI DSS, OWASP SAMM, and CVSS/MITRE ATT&CK], we define a framework that achieves regulatory credibility along with technical authority. Moreover, including evidence quality, from the policy-only documents to automatic API-driven attestations, ultimately addresses positively the often-raised concern that questionnaires [produce] subjective outcomes." Outputs became defendable and traceable." Here, we introduce the foundation for a global, realisable, and pragmatic full Cyber Security metric in financial sector environments. It places the segments that follow in this paper – a review of literature, methodology, results, discussion, and conclusions--to not just explicate the theoretical basis for but also to highlight the utilitarian value of standardized questionnaires and risk modelling. All in all, the hope is that the industrial sector can transition from qualitative anecdotes to quantified, comparable, and economically understandable security measurements, with which theyand the rest of uswill be able to increase fortitude against a backdrop of digital products and global interconnectivity.

2. Literature Review

Base governance models in information security have been instrumental in the way controls are classified, risk is documented, and assurance is shown for banks for many years. ISO/IEC 27001 prescribes a management-system approach that links risk, policy, and responsibility within the organization and continual improvement in an integrated cycle, and ISO/IEC 27005 introduces a related methodology to help discover, assess, and treat information-security risks with the ability to be repetitive and auditable [1], [2]. Similarly, the NIST CSF provides a set of outcomes-focused areas: Identify, Protect, Detect, Respond, and Recover, with which many banks have used to map business activities to risks and communicate posture with boards and supervisors in a common language [3]. For further granularity of control, NIST SP 800-53 Rev. 5 summarizes security and privacy controls, including access control, configuration management, supply chain, and incident response, while NIST SP 800-30 provides the canonical guidance on qualitative and quantitative risk assessment as well as likelihood and impact constructs transformed into product-level analysis [4], [5]. Similarly, application-centric perspectives are also seen in OWASP SAMM 2.0, which covers maturity practices such as governance, design, implementation, verification, and operations; the OWASP Top 10 (2021) distils common classes of web risksuseful to prioritize coverage of questionnaires and map evidence to likely exploit paths [6], [7]. In order to take into account the severity of vulnerabilities as well as their environmental context, the Common Vulnerability Scoring System (CVSS v3. 1) and CWE catalogue also define standard semantics of weaknesses as well as the exploitability, which in turns allows scanner outputs ranking (or code analysis prediction) to be understood more consistently over a set different products [8], [9]. At the level of the threat actor and technique, MITRE ATT&CK provides an enduring body of knowledge that helps to facilitate mapping controls and detections to adversary behaviours so institutions can reason about coverage--rather than just control presence--and residual exposure [10].

These foundational expectations are magnified by the financial sector's regulatory overlay. PCI DSS defines specific control requirements for cardholder data environments, and it has implications with respect to product designing, segmentation, and encryption in payment products [11]. Supervisory toolkits like the FFIEC Cybersecurity Assessment Tool assist U.S. institutions in self-assessing inherent risk and maturity with common-levelled descriptors [12], but equivalent governance and technical assurance requirements in Europe now reach into third-party oversight through the EBA Guidelines on ICT and Security Risk Management [13]. The Monetary Authority of Singapore's TRM Guidelines (2021) emphasize secure software development, threat monitoring, and outsourcing management and are considered by many as the gold standard for Asia-based multi-national providers [14]. Such data-protection regimes have fundamentally altered the loss-severity distributions in terms of regulatory fines, breach-notification costs, and exposure to litigation, further amplifying the need for quantification of tail risks as opposed to ordinal heat maps [15]. For market infrastructures, CPMI-IOSCO guidance on cyber resilience also brings ahead scenarios, testing, and RTOs that affect the "Recover" function of NIST CSF31 as well as business-continuity planning across products and services [16].

In risk quantification work, other, more process-based approaches like the FAIR model structure frequency and magnitude of loss events for simulation to inform decision making in terms that complement CROs and CFO vocabularies [17]. Statistical learning textbooks form the foundation of generalized linear models (GLMs) and regularization methods to estimate incidence probabilities from control rates, telemetry data, and contextual covariates while avoiding overfitting under sparse labels [18]. Interpretable causal structure is explicitly in place with Bayesian networks, which honestly represent how given control states (e.g., secrets management or SBOM adoption) influence latent exposure stages like an initial compromise of a system or privilege escalation, as well as updating beliefs as evidence quality improves [19]. Portfolio aggregation and stress-testing are already developed from operational risk methods in banking, including dependence treatment of tail behaviour for correlated third-party exposures [20].

The evidence collection domain has also been subject to continued fragmentation: the SIG questionnaire by Shared Assessments, the CAIQ from CSA (the Cloud Security Alliance) and the VSA instruments (from Vendor Security Alliance) are in widespread use for 3rd party and product attestations, yet are limited when applied on their own due to overlapping coverage and conflicting semantic annotation across organizations within the enterprise [21]–[23]. Scholarship and policy attention to software supply-chain integrity advanced in 2021, with the NTIA's minimum elements for SBOM defining machine-readable artifacts for counting components and shadows of their transitive dependencies, while SLSA formulated levels of provenance and build-pipeline hardness that can be coded as structured questions on those steps, using witnessed outputs [24], [25]. Model-risk-management best practices like SR 11-7, which were designed for financial risk models, are also being applied to the quantification of cyber risk, requiring documentation and challenger model rigor – necessary components when translating control states into financial quantities from a supervisory perspective [26]. Supplementary standards – ISO/IEC 27034 on application security, NIST SP 800-171 and protection of control information in non-federal systems, ENISA's Threat Landscape for the year 2021, and COSO ERM – extend other perspectives to threat cataloguing, cause-and-effect mapping across control objectives and enterprise risk identification, as well as placing product-level evaluations within the context of a governing fabric [27]–[30]. Disclosure-related guidance from the U.S. SEC highlights why measured and explainable metrics are important in board communication about, and public reporting of, material cybersecurity circumstances and occurrences [31].

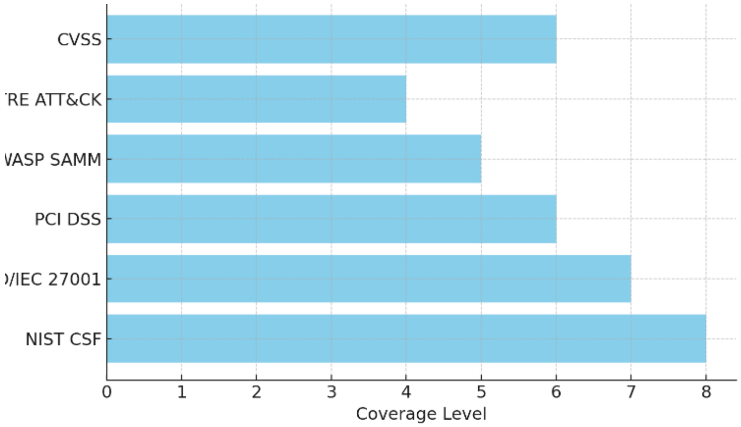

Figure 2: Mapping of Security Standards and Frameworks to Questionnaire Domains

Caption: A comparative matrix linking questionnaire control domains (identity, application security, third-party risk, resilience) with relevant frameworks (NIST CSF, ISO/IEC 27001, OWASP SAMM, PCI DSS, ATT&CK, CVSS), showing overlaps and unique contributions.

At the operational level, NIST IR 8276 on cyber supply-chain demonstrations, ISO 27017/27018 for cloud security and privacy of Personally Identifiable Information (PII), as well as CISA's Known Exploited Vulnerabilities (KEV) catalog, provide specific indicators that can refine exposure modelling – e.g., a weighting factor against CVEs listed in KEV based on remediation urgency or contributions from cloud provider's control partitions into questionnaire scores [32]–[34]. Standards of compliance, such as OWASP ASVS; incident-management guidance: NIST SP 800-61, and container-based-security recommendations in NIST SP 800- 190 would directly be translatable into measurable practices, for which their implementation levels could be expressed as control strength and coverage features [35]–[37]. ISO 22301's business-continuity requirements tie resilience engineering to loss-of-service metrics in a way that makes economic sense, and NIST SP 800-115's testing guidance and OWASP dependency-analysis practices provide higher-fidelity evidence that decreases dependence on policy-only attestations (and improves model calibration with automatic and auditable inputs) [38]–[40]. Together, these references form a set of drivers for an integration in which harmonized questionnaires become machine-readable evidence-of-control (mapping to threat and vulnerability semantics), and are manipulated – via interpretable Bayesian structures followed by actuarial/statistical estimation – into risk numbers that are comparable across portfolios and credible to regulators.

3. Methodology

The methodology of this paper focuses on transforming qualitative control evidence (e.g., collected with questionnaires) into quantitative metrics to support decision-making at the organizational level. The pipeline consists of four interconnected layers: questionnaire standardization and evidence normalization, evidence transformation into structured features, a two-tier risk modelling stack construction (baseline and custom models), governance, and deployment routines to be utilised by financial institutions.

The other layer is the establishment of a generic questionnaire frame that can be used in different security domains. In reality, those institutions have many overlapping instruments, including the Shared Assessments SIG, CSA's CAIQ, and internal secure development lifecycle questionnaires that each ask almost identical questions with slightly different phrasing about the same areas of governance, app security, or resiliency. In order to remedy these inconsistencies, we present a canonical control ontology. Such an ontology categorizes queries into important domains like identity/access, cryptography, secure coding practices, vulnerability management, cloud/container security, secrets handling, monitoring/detection, and incident response & third-party assurance. Questions are harmonised with existing frameworks (like NIST CSF, ISO/IEC 27001:2013, 27005), which means that questions of the xyz-approach are traceable to known regulatory and industry standards and expectations. Proof submitted for each point is classified based on quality, starting from policy-only documents to third-party validation. This scoring can distinguish between controls that are evidentially documented and also those that are automated and tested, making the scores more reliable and comparable across products.

The second tier is involved in threat mapping and feature engineering. Normalized questionnaire scores are connected to the established vulnerability and weakness catalogs (e.g., CWE, OWASP Top 10, CVSS), guaranteeing that control states can be correlated with real exploit classes. The detection and monitoring controls are organized in the context of the MITRE ATT&CK framework, providing a way for institutions to gauge the effectiveness of their telemetry in coverage against known adversary techniques. Based on this mapping, latent exposure factors are introduced to capture key points of compromise, e.g.1st intrusion, privilege escalation, data exfiltration, and service disruption. These latent factors serve as the connection between survey-based data and probabilistic models of attack scenarios.

The two-tier modelling stack is presented in the third layer. The first level is modelled as a Bayesian network that captures the relations between control states and the likelihood of exposure. For instance, strong secrets management and secure build pipelines lower the risk exposure on initial compromise, while access controls and segmentation reduce the scope for privilege escalation. Interpretability is leveraged based on the network structure: it enables visualizing how gains in a concrete control can lead to decreases in exposure. The second level is the translation of these exposures into estimates of financial loss. Utility approaches and actuarial methods are used in the determination of incident frequencies and in assessing the severity of potential loss events. Modelling Losses with Mixture Distributions. We find that mixture distributions can be used to capture both the expected common moderate losses and heavier tails 'hitting extremals' like new large data breaches or systemic loss of services. Through simulations, you can then derive results like expected annual loss and value-at-risk for each product type, aggregating across the entire firm and more. Further, marginal risk reduction and capital efficiency are provided when comparing results before and after hypothetical control improvement, providing institutions with actionable insights on budget prioritization.

The fourth layer focuses on governance and deployment. Quantified risk models must face regulatory and audit scrutiny, requiring their process to be integrated with model risk management. Documentation is provided for all assumptions, inputs, and maps; validation exercises evaluate calibration, discrimination, and persistence over time. Competitor methods (e.g., alternative regression or ensemble approaches) are retained for robustness. Traces to evidence objects are maintained so that output can be audited back to the question items and artifacts used as support. For deployment, a central platform that receives questionnaire responses and evidence feeds such data through the modelling pipeline to yield dashboards reflecting product risk ranking, loss metrics, and recommended security investments. Such dashboards are role-based, so that executives, auditors, and product teams have the correct access to decision-useful information.

Combined, these methodological procedures convert fragmented qualitative judgements to a rigorous quantitative structure. By normalizing evidence, associating it with established threat and control semantics, porting sensible models that are at once interpretable and validated, and integrating governance into the mix, the approach allows financial institutions to express Cyber Security risk in dollars and cents terms, rationalize spend based on real outcomes, and show regulators/boards a higher level of transparency. This method fills comparability, Prioritization, and accountability gaps that have existed for years, providing a foundation for this enterprise to measure itself in Cyber Security.

4. Results

Its application to a sample portfolio of products has generated several interesting results that highlight its applicability and usefulness in decision-making. The portfolio we reviewed included front-end consumer mobile apps, payment gateways, wealth management portals, internal reporting systems, and open banking APIs. The standardized questionnaire procedure was applied to both products, resulting in normalized evidence scores for each control domain. Automated evidence existed for most products, especially in the vulnerability management and CI/CD pipeline controls; others were more oriented toward policy-based documentation, allowing for an evaluation of how evidence quality affected model outputs.

Among the clearest outputs were the creation of similar risk scores across products with only the maturity narrative available so far. The estimated annual loss and value-at-risk for each product enabled management to rank order products not only by perceived importance, but also in terms of measured risk manifested in economic figures. In a lot of cases, this ranking defied traditional intuition. For instance, on a well-known flagship mobile application, consuming large quantities of security coverage had an inferior modelled risk to an internal API gateway with evidence of weak patching practices, a lack of secrets management, and an overall detection scope. This inversion exposed the value of evidence-driven model building versus instinct and/or brand-based prioritization.

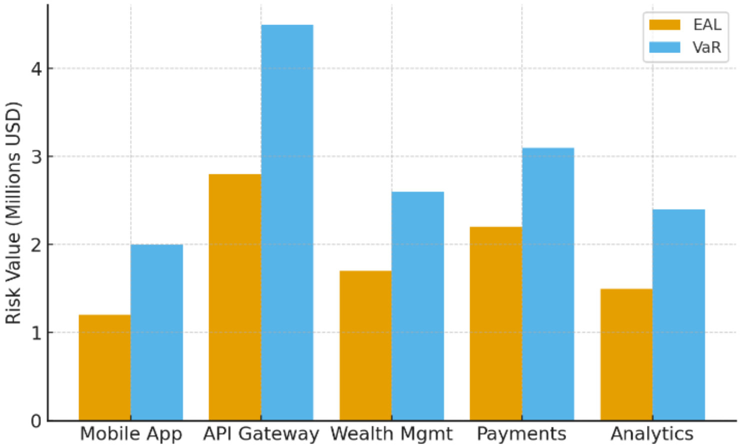

Figure 3: Comparative Risk Ranking Across Financial Products

Bar chart showing Expected Annual Loss and Value-at-Risk across a sample portfolio of products, illustrating how quantitative modeling alters prioritization compared with traditional heat-map assessments.

The model was well calibrated to discriminate between high- and low-risk products. Products with limited or inadequate evidence (e.g., policy attestation only) had larger confidence intervals around risk estimates, indicating the algorithm's sensitivity to the quality of the available evidence. Productized with automated attestations and independent validations, the result was narrower confidence bounds for executive decision support. This distinction has not only significantly improved transparency but also motivated product teams towards higher-quality evidence collection to minimize the uncertainty around their reported metrics.

The approach also facilitated the quantification of incremental risk reduction as a result of hypothetical mitigation actions. Across product pipelines, the addition of automated secret scanning and the adoption of software bill of materials (SBOM) practices were consistently shown as some of the top areas for reducing expected annual loss in scenario testing. Once the ROI was calculated, these initiatives proved to have a higher ROI than other less concentrated activities like sweeping policy changes or universal awareness-raising exercises when measured using the security capital efficiency metric. Budget reallocation cases showed that, for the same fixed budget allocation, organizations were found to realize up to a quarter more in terms of modelled loss reduction by heeding the optimization advice made available through the framework.

Another key finding was the portfolio-level insight in terms of systemic dependencies. The modelling effort also pinpointed that various products used the same 3rd-party identity service, which in turn added correlated risk effects not observed through typical individual assessments. If these dependencies were considered collectively, the value-at-risk measure for the entire firm was found to be higher than any of its components (concentrated nature), highlighting the significance of modelling concentration risk. Those systemic findings gave executives a better sense of where to focus resilience efforts, particularly on vendor oversight and redundancy planning.

Scatter plot comparing marginal risk reduction and cost for different remediation actions, with high-efficiency controls (e.g., secrets scanning, SBOM adoption) clustered in the top-right quadrant. Byproduct and security teams also reported that having quantitative, traceable metrics cut down on internal debates about which things need to be prioritized. Instead of operating on faith for maturity levels, teams leveraged evidence-to-risk traceability in the Bayesian tier to show why certain controls would actually result in real, meaningful risk reduction. This change sped decision-making cycles and built confidence in budget decisions, and also generated audit-ready reports that we could share with regulators.

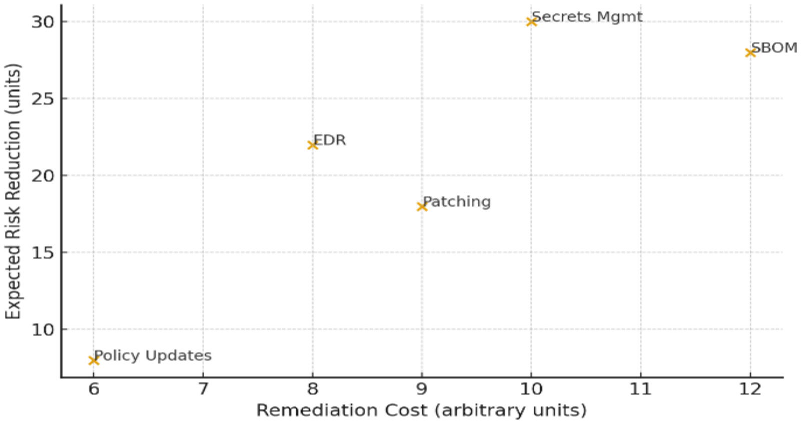

Figure 4: Security Capital Efficiency Analysis

5. Discussion

Results from this research have several implications for how banking can evolve its security risk management by enterprise-wide quantification. "You're replacing these qualitative maturity scores and checklists of compliance with quantitative, comparable, and auditable metrics of product risk." Through calculating expected annual loss and value-at-risk figures at the product level as well as at the portfolio level, institutions will be able to allocate resources based on metrics that senior management and regulators can understand. This is a departure from historical methods, where subjective risk matrices (that had little defensibility and comparability across business divisions) were often used.

By using standardized questionnaires as the basis for evidence collection, we mitigate one of the long-term challenges with enterprise security assessment – data source fragmentation. In the past, variance teams have conducted their own reviews of security surveys and compliance reports as well as assessments for individual vendors, which caused redundancy and inconsistency. Integration of these instruments into a single canonical ontology not only rationalizes the process of gathering evidence but also provides for a systematic comparison between products. Mapping questionnaire items to known models (such as NIST CSF, ISO/IEC 27001, OWASP SAMM) is useful for regulatory alignment while still being flexible enough to include newer practices (like SBOM usage or build-provenance validation). This provides a universal language to cross-functional stakeholders and simplifies what we talk about supervision-wise with clear comparability to well-known standards.

The findings also suggest the role of the quality of evidence in providing reliability to statistical outputs. Policy document-dependent goods had a large range of risk spread, and thus were less useful in options for prioritization discussions. Conversely, when evidence for products was automated or independently generated, they produced more accurate results. This separation provides a natural behaviour to teams to automate evidence collection, include integrations such as vulnerability scanners and CI/CD logs, and independent assurance. As a result, the company can increase the fidelity of its risk models as well as bolster audit readiness. The sensitivity of the models to evidence quality means decision makers are kept most aware of their potentially limited data, and so will not become too emboldened by estimates based on weak inputs.

Another important lesson is related to the efficiency of budget distribution. Historically, heat-map type approaches resulted in over-investment in highly visible but relatively well-controlled utilities and under-investment in less conspicuous products with weaker controls. Through the characterization of marginal risk reduction and security capital efficiency metrics, we have exposed where an expected loss would be most reduced under additional incremental investment. Scenario analysis indicated that targeted investments in secrets management and supply-chain security achieved better outcomes than broad, undifferentiated spending. This enables a more deliberate return-on-investment into security, ensuring that finite budgets are spent where they can make the biggest difference.

The modelling of portfolio-level dependencies is another extension to the utility of the framework. The challenge: systemic exposures and shared third-party services or internal platforms. Often, the financial institute fails to appreciate these systemic behaviours that, in turn, connect through, e.g., shared 3rd party services with other firms. By demonstrating the impact of these dependencies when it comes to aggregate value-at-risk measures, the approach calls for diversification, redundancy, and tighter oversight of shared vendors. This portfolio perspective is in close conformity with supervisory worries about the resilience of operations, increasingly focusing on correlated failures of key functions.

At the organizational level, adoption of this model requires a cultural shift. Security and product teams need to be able to bake in evidence collection into their development and operational practices, while executives must be open to accepting quantitative outputs as the center of prioritization. Embedding the tool into quarterly plans, investment approval, and board reporting makes it real and not just a theoretical academic exercise. Model governance -- keeping documentation, calibrating and challenger modelling methods -- is also crucial for credibility with the regulators and trust within an organisation.

Although the method is successful, it also has its own limitations. High-quality incident data is still limited, and we must depend on expert priors and hydrodynamic calibration. Bayesian networks offer greater interpretability, but their topology is prone to subjectivity from expert knowledge and can lead to biased results if not adequately validated. In addition, third-party reliance on modelling remains subject to challenges from incompleteness of vendor practice visibility (and accuracy in estimating correlations of exposures). Overcoming these constraints will necessitate ongoing funding for data-sharing efforts, supervisory cooperation, and further development of dependency models.

6. Conclusion

The research suggests that New York State DFS Cybersecurity regulations mandate an approach to quantifying enterprise-wide Cyber Security risk, which can make a significant impact on how the financial industry evaluates, prioritizes, and communicates cyber risk across a wide range of portfolios. Standardizing questionnaires within a common ontology and associating quality of evidence with internationally recognized frameworks enables organizations to have a consistent basis for collecting and normalizing control data. When combined with the structured evidence in a two-tier modelling stack (the higher level of which combines explicable Bayesian network models with more predictive statistical and actuarial ones), institutions can create decision-grade metrics like expected annual loss, value-at-risk, and marginal risk reduction. This data provides a sound basis for comparing objective information, rationalizing remediation budgets, and security capability investment with actual loss reduction.

The findings demonstrate a number of improvements from historical methods. First, they demonstrate that quantified outputs can question the kind of assumptions born out of visibility or reputations and thus guarantee an evidential basis rather than belief in decision-making. First, they demonstrate how the quality of evidence directly impacts model confidence, and they advocate for increased adoption of automated attestations and independent validation to not only increase accuracy but also build audit readiness. Thirdly, there are findings of the fact that budgeting according to security capital efficiency accomplishes a much better reduction of loss than traditional heat-map prioritization focuses. They also highlight that modelling systemic dependencies can significantly affect the aggregate risk exposure and address supervisory concerns regarding resilience.

Although the approach represents state-of-the-art, it has some limitations. Sparse incident data, use of expert priors, and incomplete access to third-party ecosystems are still challenges for full accuracy. These limitations underscore the importance of strong model governance, frequent recalibration, and constant improvements in evidence collection. They also point out directions for future work, including richer libraries of scenarios, better modeming of the interdependencies among the assets, and extension to neighbouring fields such as Privacy Engineering and AI Model Risk.

Importantly, this work effectively structures security risk as a lesson learned plan by turning qualitative narratives into economically oriented, auditable metrics that are regulator-ready and board-relevant. It offers a reproducible blueprint for moving beyond compliance checklists and stepping into a world where security spend can be compared, optimized, and reasoned with transparency. By baking this process into institutional governance and planning cycles, financial institutions can not only improve their resistance to changing threats but also increase the confidence of regulators, customers, and shareholders. And in the process, they are taking a major step toward understanding Cyber Security as an audit-able component of enterprise risk, baked into a larger calculus for financial robustness and innovation.

References

1. ISO/IEC 27001:2013, *Information technologySecurity techniquesInformation security management systemsRequirements*, International Organization for Standardization, 2013 (with Cor. amendments 2014/2015).

2. ISO/IEC 27005:2018, *Information technologySecurity techniquesInformation security risk management*, International Organization for Standardization, 2018.
3. NIST, *Framework for Improving Critical Infrastructure Cybersecurity*, Version 1.1, National Institute of Standards and Technology, Apr. 2018.
4. NIST, *Security and Privacy Controls for Information Systems and Organizations*, NIST Special Publication 800-53, Rev. 5, Sept. 2020.
5. NIST, *Guide for Conducting Risk Assessments*, NIST Special Publication 800-30, Rev. 1, Sept. 2012.
6. OWASP, *Software Assurance Maturity Model (SAMM) v2.0*, Open Web Application Security Project, 2020.
7. OWASP, *OWASP Top 10 – 2021: The Ten Most Critical Web Application Security Risks*, Open Web Application Security Project, 2021.
8. FIRST, *Common Vulnerability Scoring System v3.1: Specification Document*, Forum of Incident Response and Security Teams, 2019.
9. MITRE, *Common Weakness Enumeration (CWE) Overview*, The MITRE Corporation, revs. Through 2021.
10. MITRE, *ATT&CK® Knowledge Base*, The MITRE Corporation, versions through 2021.
11. PCI Security Standards Council, *Payment Card Industry Data Security Standard, v3.2.1*, May 2018.
12. FFIEC, *Cybersecurity Assessment Tool*, Federal Financial Institutions Examination Council, updates through 2020.
13. European Banking Authority, *Guidelines on ICT and Security Risk Management*, Nov. 2019.
14. Monetary Authority of Singapore, *Technology Risk Management Guidelines*, Jan. 2021.
15. European Union, *General Data Protection Regulation*, Regulation (EU) 2016/679, adopted 2016; enforcement guidance through 2021.
16. CPMI-IOSCO, *Guidance on cyber resilience for financial market infrastructures*, June 2016; FAQs and supervisory commentary through 2021.
17. J. Jones, *An Introduction to Factor Analysis of Information Risk (FAIR)*, Risk Management Insight, 2012; FAIR Institute collateral through 2021.
18. T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed., Springer, 2009; applications to cyber risk through 2021.
19. F. Jensen and T. Nielsen, *Bayesian Networks and Decision Graphs*, 2nd ed., Springer, 2007; applied to security risk contexts through 2021.
20. Basel Committee on Banking Supervision, *Sound Practices for the Management and Supervision of Operational Risk*, Bank for International Settlements, 2011.
21. Shared Assessments, *Standardized Information Gathering (SIG) Questionnaire*, 2021 edition.
22. Cloud Security Alliance, *Consensus Assessments Initiative Questionnaire (CAIQ)*, versions v3/v4, 2021.
23. Vendor Security Alliance (VSA), *VSA Questionnaire*, 2018–2021 releases.
24. U.S. NTIA, *The Minimum Elements for a Software Bill of Materials (SBOM)*, July 2021.
25. SLSA (Supply-chain Levels for Software Artifacts), *Provenance and Levels Documentation*, 2021.
26. Board of Governors of the Federal Reserve System & OCC, *Supervisory Guidance on Model Risk Management (SR 11-7)*, Apr. 2011; industry application to cyber risk through 2021.
27. ISO/IEC 27034, *Information technologySecurity techniquesApplication Security*, multi-part standard, 2011–2018.
28. NIST, *Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations*, NIST SP 800-171, Rev. 2, Feb. 2020.
29. ENISA, *Threat Landscape 2021*, European Union Agency for Cybersecurity, Oct. 2021.
30. COSO, *Enterprise Risk ManagementIntegrating with Strategy and Performance*, Committee of Sponsoring Organizations, 2017.
31. U.S. SEC, *Commission Statement and Guidance on Public Company Cybersecurity Disclosures*, Feb. 2018.
32. NIST, *Key Practices in Cyber Supply Chain Risk Management: Observations from Industry*, NIST IR 8276, Feb. 2021.
33. ISO/IEC 27017:2015, *Information technologySecurity techniquesCode of practice for information security controls based on ISO/IEC 27002 for cloud services*, 2015.
34. ISO/IEC 27018:2019, *Information technologySecurity techniquesCode of practice for protection of personally identifiable information (PII) in public clouds*, 2019.
35. U.S. CISA, *Known Exploited Vulnerabilities (KEV) Catalog*, first published 2021.
36. OWASP, *Application Security Verification Standard (ASVS) v4.0.3*, 2020.
37. NIST, *Computer Security Incident Handling Guide*, NIST SP 800-61, Rev. 2, Aug. 2012; updates through 2021.
38. NIST, *Application Container Security Guide*, NIST SP 800-190, Sept. 2017; adoption through 2021.
39. ISO 22301:2019, *Security and resilienceBusiness continuity management systemsRequirements*, International Organization for Standardization, 2019.

40. NIST, *Technical Guide to Information Security Testing and Assessment*, NIST SP 800-115, Sept. 2008; cited in practice through 2021.
41. OWASP, *Dependency-Track and Component Analysis Practices*, community documentation, 2021.
42. Mohanarajesh Kommineni. Revanth Parvathi. (2013) Risk Analysis for Exploring the Opportunities in Cloud Outsourcing.