iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

CYBER RISK QUANTIFICATION UNDER UNCERTAINTY: INTEGRATING THREAT PROBABILITY, BUSINESS IMPACT, AND RECOVERY COSTS

Mojisola Ologe
Coauthor: Aisha Abdullahi

ABSTRACT
Digital infrastructure, connected information systems and third-party technology service providers are vital to the business operations of organisations. While these relationships can bring efficiencies and expansion, they also present a risk for cybersecurity incidents that could lead to operational disruption; financial loss; regulatory implications; and a significant recovery cost. While traditional qualitative risk assessments can help prioritize risks initially, they often do not adequately capture the uncertainty, financial impacts, and interdependencies of cyber incidents, as they use a broad classification of low, medium and high risk. In this paper, a conceptual quantitative approach to combining threat-event frequency, conditional compromise probability, operational disruption, financial exposure, and recovery costs into a single cyber-risk assessment model is developed. The framework is based on well-known risk-management principles and on the economics of Cyber security, considering incident frequency and severity in a probabilistic way, separating direct losses and operational and recovery costs and including parameter uncertainty and dependencies among the risk scenarios. Loss exceedance probabilities and conditional value at risk are offered as additional decision metrics in addition to the expected annual loss, where extreme losses are material. The framework also ties residual risk assessments to decisions on investing in cybersecurity within a constrained budget. A procedure for model development is proposed, which involves parameter estimation, simulation, sensitivity analysis, and verification against the baseline approaches, all performed in a quantitative manner. Since no measured improvements in prediction or investment efficiency are claimed, the model has not yet been empirically implemented and validated. Instead, it offers a structured basis for future empirical testing as well as for more visible and evidence-based cyber security investments and risk-prioritization decisions.

Keywords:
cyber risk quantification, cybersecurity investment, threat probability, operational disruption, recovery costs, uncertainty analysis, expected annual loss, risk prioritization

I. INTRODUCTION

1.1 Background to the study
Given the reliance of business operations on information systems, digital infrastructure, cloud services, and third-party providers that are increasingly interconnected in today's business environment, cybersecurity has emerged as an important part of organizational risk management. An attack on one of these resources can impact not just the confidentiality, integrity, or availability of information, but more. It can halt production, slow down service delivery, raise operating costs, provide access to sensitive data, create legal or regulatory requirements, and reduce customer trust. Cyber risk is thus a technical and economic issue. Anderson and Moore (2006) suggest that the economic incentives and distribution of security costs and benefits drive information-security outcomes. Gordon and Loeb (2002) also show how the decision about cybersecurity investments can be viewed as a relationship between security spending and the expected loss. These concepts highlight the need to have more than a list of technical weaknesses – they need a way to link up security incidents with the consequences to their business.
An important starting point can be risk assessment practices. The National Institute of Standards and Technology (NIST, 2012) suggests that the following should be identified: threat sources and events, vulnerabilities, likelihood, and adverse impacts. The Open Group (2020) offers a quantitative measure to quantify information risk, and ISO/IEC 27005:2022 offers guidance for information-security risk management. But in practice making investment decisions based on these principles is challenging as incident data are scarce, business impact is unpredictable, and recovery costs can differ significantly from incident to incident.
For instance, the loss of two organizations attacked by ransomware might be different. You might have good backups, recovery plans that you have practiced and other arrangements for service. The other could depend on a critical system and not have a tested recovery plan. If their exposure is equal, their operational disruption, financial costs during recovery and overall financial impact can be significantly different. This separation may underscore

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

the importance of a model that not only calculates the probability of occurrence of a security event, but also the probability of successfully compromising the business, the magnitude of that business impact, and the costs to recover from the impact. It should also reflect uncertainty in the underlying estimates and how alternative security control options might affect the organisation's residual exposure.

1.2 Statement of the problem

It is a common occurrence in organizations to have to make cybersecurity investments while lacking full knowledge of the severity, frequency and financial impact of potential future attacks. There may be insufficient or conflicting historical data, especially for rare but extreme events. Cyber risk research is hindered by data availability, identified by Cremer et al. (2022), and costs and causes of reported cyber incidents are analyzed by Romanosky (2016).

The following are three problems that are directly related to these limitations.

First, the probability of a threat is not easily estimated when there are few incidents, threat conditions change and past controls affect the likelihood of a threat becoming a successful incident. A vulnerability count or threat rating is not necessarily a prediction of financial impact.

Secondly, the impact of business is a multi-dimensional phenomenon. An incident could result in loss of income, extra labor cost, contract penalties, customer compensation, investigation expenses, and future ramifications. Assessments may fail to consider the losses if these categories are not clearly defined, or the same loss can be counted multiple times.

Thirdly, recovery costs are not always viewed as a separate part of the cyber risk. The cost of incident investigation, system restoration, specialist assistance, data reconstruction, business continuity and post incident improvement may be significant. The size of these depends on the characteristics of the incident, organizational readiness and the length and complexity of the recovery.

This means that decision-makers can get an inadequate level of transparency and/or may not be able to compare risk estimates on a scenario-by-scenario basis. A threat-event frequency, threat severity, business disruption, financial exposure and recovery cost model may serve as a more consistent basis for comparing risks and evaluating alternative controls.

1.3 Aim and objectives

The aim of this paper is to develop a conceptual quantitative cyber-risk model that integrates threat probability, business impact, and recovery costs under uncertainty to support cybersecurity investment and risk-prioritization decisions.

The specific objectives are to:

1) Examine existing cyber-risk assessment and cybersecurity investment approaches.
2) Define the principal variables required to estimate the frequency and consequences of cyber incidents.
3) Develop a mathematical framework for integrating operational disruption, financial exposure, and recovery costs while avoiding double-counting.
4) Incorporate parameter uncertainty and dependencies between incident scenarios.
5) Establish a simulation-based methodology for estimating aggregate losses and comparing risk-reduction controls.
6) Propose criteria for evaluating the model against simpler risk-assessment approaches.

1.4 Research questions

The study addresses the following questions:

1) How can organizations estimate cyber-incident frequency and conditional compromise probability when historical data are incomplete?
2) How can operational disruption, financial exposure, and recovery costs be integrated into a common quantitative model without double-counting?
3) How does uncertainty in incident frequency, severity, and recovery duration affect estimated cyber risk?
4) How can the model support the comparison and prioritization of cybersecurity investments under budget constraints?
5) What evaluation procedures are necessary to determine whether the proposed framework provides useful decision support beyond simpler approaches?

1.5 Research propositions

Because the framework is conceptual and has not yet been empirically tested, the following propositions are presented for future investigation rather than as established findings.

- P1: Integrating operational disruption, direct financial exposure, and recovery costs will produce a more comprehensive estimate of cyber loss than an approach that considers only one or two of these components.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

- P2: Explicitly representing uncertainty will reveal a wider range of plausible losses than reporting point estimates alone.
- P3: Accounting for dependencies between cyber scenarios will materially affect aggregate-loss estimates where incidents share infrastructure, providers, or critical systems.
- P4: Investment rankings based on residual-loss distributions may differ from rankings based solely on qualitative risk scores or expected annual loss.

1.6 Significance and scope

The paper contributes a structured model-development approach that links cyber-risk assessment to business consequences and investment decisions. Its potential value lies in making assumptions visible, separating distinct loss categories, representing uncertainty, and comparing the effects of controls across plausible incident scenarios. The framework is intended for organizations seeking to prioritize cybersecurity expenditure, particularly where resources are constrained and business leaders need an interpretable account of residual risk. It is not limited to a particular industry, although implementation would require sector-specific data and assumptions.

The study is limited to conceptual model development and proposed quantitative evaluation. It does not claim to have collected original organizational data, run a completed simulation, or demonstrated empirically that the framework outperforms existing methods.

II. LITERATURE REVIEW

2.1 Conceptual foundations of cyber-risk quantification

Cyber risk concerns the possibility that a cybersecurity-related event will produce adverse consequences for an organization's assets, operations, objectives, or stakeholders. Its assessment requires consideration of threat sources, vulnerabilities, existing controls, the likelihood of adverse events, and the severity of their consequences. NIST (2012) provides a structured method for conducting such assessments. Quantification extends this process by representing relevant variables numerically. Rather than describing ransomware risk as simply “high,” a quantitative assessment might estimate the annual frequency of ransomware incidents, the distribution of service-interruption durations, and the financial consequences of different recovery outcomes. Such estimates allow decision-makers to compare scenarios and evaluate how alternative controls may change residual exposure.

However, quantitative representation does not eliminate uncertainty. The quality of an estimate depends on the underlying evidence, the assumptions used, and the suitability of the model. Hubbard and Seiersen (2016) argue for more rigorous measurement of cybersecurity risk, including the explicit treatment of uncertainty. The Open Group (2020) similarly provides a structured approach to quantitative risk analysis through the Factor Analysis of Information Risk (FAIR) methodology.

The proposed framework builds on these foundations by combining event frequency and severity with distinct business-loss categories and recovery expenditure. Its contribution is not the claim that quantitative cyber-risk assessment is new, but the specification of an integrated decision-support structure that makes loss composition, dependencies, uncertainty, and investment implications explicit.

2.2 Threat probability and incident frequency

Threat probability is an essential input to cyber-risk estimation, but it must be defined carefully. The probability that a threat actor attempts an attack is not necessarily the same as the probability that the attack succeeds, causes a business incident, or produces a financial loss.

A useful assessment therefore distinguishes at least three elements:

1) Threat-event frequency: how often a relevant threat event or attack attempt is expected to occur.
2) Conditional compromise probability: the likelihood that an attempt succeeds, given the threat event and the organization's vulnerabilities and controls.
3) Loss-producing incident probability: the likelihood that a successful compromise results in a defined adverse business outcome.

These quantities may be estimated from incident histories, vulnerability and control assessments, threat intelligence, structured expert elicitation, and external studies. However, evidence from another organization may not transfer directly because exposure, architecture, attacker interest, control maturity, and reporting practices differ.

NIST (2012) emphasizes the need to assess threat events in relation to vulnerabilities, predisposing conditions, and existing safeguards. The Open Group (2020) also provides a quantitative structure for estimating threat-event frequency and vulnerability. Such approaches support scenario-specific assessment rather than relying solely on generic threat rankings.

Where sufficient incident-count data exist, annual incident frequency may be modelled using a Poisson distribution. This approach assumes that events occur independently at a stable average rate over the period. Those

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

assumptions may not hold when threat activity changes over time, incidents cluster, or one event triggers multiple related failures. Alternative count models or time-varying rates may therefore be necessary.

2.3 Operational disruption and business impact

Operational disruption represents the consequences of an incident for an organization's ability to deliver products or services. Examples include unavailable business applications, halted production, inaccessible customer records, delayed transactions, and reduced staff productivity.

The financial consequences of disruption depend on the affected service, the duration of the interruption, the number of customers or processes affected, and the availability of alternative arrangements. An outage lasting two hours may be manageable for one organization but severe for another whose operations depend on continuous access to a particular system.

A practical model should therefore connect technical events to specific business services. This requires identifying critical systems, mapping dependencies, estimating service-recovery times, and determining how interruption affects revenue, costs, and contractual obligations. NIST's contingency-planning guidance (Swanson et al., 2010) provides a basis for considering system recovery, operational continuity, and restoration requirements.

Operational impact should not be equated automatically with lost revenue. Revenue delayed rather than permanently lost, costs avoided during an outage, substitute-service expenditure, and lost production may have different financial implications. A model must specify how each is valued to avoid overstating the consequences.

2.4 Financial exposure and recovery costs

Financial exposure may include direct response expenditure, lost contribution from disrupted operations, customer compensation, legal and regulatory costs, contractual penalties, and other incident-related losses. Romanosky (2016) examines the costs and causes of cyber incidents, illustrating the importance of looking beyond the technical event itself. Cavusoglu et al. (2004) also found that public announcements of Internet security breaches were associated with negative market reactions for affected firms. Such findings indicate that incident consequences may extend beyond immediate technical remediation, although market-value changes should not be added automatically to other financial loss estimates.

Recovery costs should be defined as the expenditure required to investigate, contain, remove, restore, and validate affected systems and services, together with other explicitly specified recovery activities. Cichonski et al. (2012) provide guidance on incident handling, while Swanson et al. (2010) discuss contingency planning and the restoration of information systems.

The distinction between operational losses and recovery expenditure is important. If an organization estimates business interruption using a daily loss figure that already includes emergency labour and recovery expenses, adding those same expenses separately would double-count them. Conversely, excluding recovery costs because they are difficult to estimate would understate total exposure.

The proposed model therefore defines separate loss categories and requires the analyst to document the boundaries of each category before estimating its value.

2.5 Uncertainty in cyber-risk estimation

Cyber-risk estimates are uncertain for several reasons. First, aleatory uncertainty arises from the variability of events themselves, such as the number of incidents occurring in a year or the duration of a particular outage. Second, epistemic uncertainty arises from incomplete knowledge, such as limited evidence about attacker behaviour, control effectiveness, or the financial consequences of a rare event. Third, model uncertainty arises when the mathematical representation does not fully capture the real-world process.

These forms of uncertainty have different implications. Random variation may persist even when extensive data are available. Knowledge-related uncertainty may be reduced through better data collection, testing, or expert assessment. Model uncertainty requires alternative specifications, sensitivity analysis, and validation against observed outcomes.

The Guide to the Expression of Uncertainty in Measurement (Joint Committee for Guides in Metrology [JCGM], 2008) provides general principles for expressing uncertainty, while Saltelli et al. (2008) explain methods for examining how uncertainty in input variables affects model outputs. Although these sources are not cyber-risk-specific, their methods are relevant to quantitative risk analysis.

Cremer et al. (2022) identify limitations in cyber-risk data availability, reinforcing the need to avoid presenting estimated parameters as exact values. Where direct evidence is weak, the model should document the evidence source, uncertainty range, and rationale for each assumption.

2.6 Cybersecurity investment economics

Risk quantification becomes especially useful when it supports decisions about whether, when, and how much to invest in cybersecurity. Gordon and Loeb (2002) analyse the economic relationship between information-security

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

expenditure and expected loss. Their work provides a foundation for evaluating security spending as an investment decision rather than assuming that greater expenditure always produces proportionately greater protection.

Anderson (2001) explains why information security presents distinctive economic challenges, including incentive problems and the difficulty of aligning the interests of different parties. Fedele and Roner (2022) review research on cybersecurity investment and distinguish models involving isolated firms from those involving interdependent organizations.

Interdependence matters because an organization's risk can be influenced by external service providers, suppliers, customers, and shared infrastructure. Böhme and Schwartz (2010) discuss the distinctive features of cyber-insurance models, including correlated risk and interdependent security. These issues are relevant to a model that must account for third-party incidents and shared dependencies.

A decision framework should therefore consider both the estimated reduction in loss and the cost of implementing a control. It should also recognize that controls may reduce different components of risk. Multi-factor authentication may reduce the probability of unauthorized access; endpoint detection and response may reduce the duration or severity of a compromise; tested backups may reduce recovery time and restoration expenditure. A model that represents these effects separately can provide a clearer explanation of why one investment may be preferable to another.

2.7 Existing frameworks and the research gap

Several established approaches provide components of the proposed model. NIST (2012) offers a structured risk-assessment process. ISO/IEC 27005:2022 provides guidance on information-security risk management. The Open Group's FAIR standard (2020) offers a quantitative approach to analysing information risk. NIST publications on incident handling, contingency planning, and security controls support the operational and recovery dimensions of assessment.

These frameworks should not be treated as interchangeable. Some primarily structure risk identification and assessment; others focus on quantitative loss analysis, incident response, or system recovery. The proposed framework draws on these complementary functions rather than suggesting that existing standards ignore business impact or uncertainty.

The research opportunity is to specify a transparent model that connects these elements within a common analytical workflow: incident frequency, conditional compromise, operational impact, financial loss, recovery costs, uncertainty, dependency effects, and control investment. Its value must be evaluated empirically rather than inferred from its conceptual completeness alone.

2.8 Conceptual framework

The proposed relationship among the principal variables is shown below.

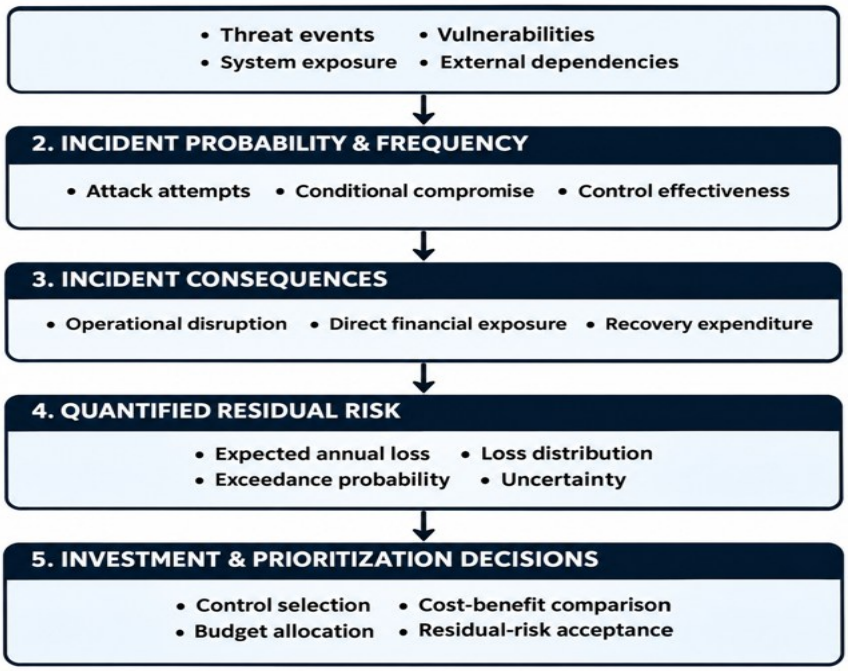


Figure 1. Conceptual framework for integrating cyber-event probability, business consequences, recovery costs, and investment decisions. Source: Author's synthesis of the cited literature.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

The framework assumes that uncertainty affects both incident likelihood and consequences. Dependencies may also connect multiple scenarios, meaning that a single external event can generate several losses. Security controls may influence incident frequency, compromise probability, loss severity, or recovery duration. The model must preserve these relationships when estimating aggregate risk.

III. RESEARCH METHODOLOGY

3.1 Research design

This study adopts a quantitative model-development design supported by a literature review and a proposed simulation-based evaluation. The purpose is to define a reproducible analytical framework rather than report results from an unperformed empirical study.

The methodology comprises five stages:

1. Define incident scenarios and their associated business services.
2. Identify variables and probability distributions representing incident frequency and severity.
3. Estimate the financial consequences of each incident, including operational disruption and recovery expenditure.
4. Simulate aggregate annual losses under baseline and alternative-control scenarios.
5. Evaluate model behaviour, sensitivity, uncertainty, and decision usefulness against simpler assessment approaches.

This design is consistent with the objective of developing a quantitative decision-support model. It also creates a clear path for future empirical implementation using organizational incident data.

3.2 Data requirements and potential sources

A practical implementation would require data from several sources. Internal records would ideally provide incident counts, incident classifications, service outages, recovery times, response expenditure, and the costs of corrective actions. Financial and operational records would support the valuation of lost production, disrupted services, and additional operating costs.

External sources could supplement internal evidence. Cremer et al. (2022) review the availability of cyber-risk data, while Romanosky (2016) examines incident costs and causes. Reports such as the European Union Agency for Cybersecurity's ENISA Threat Landscape 2023 and IBM Security's Cost of a Data Breach Report 2023 can provide contextual evidence, but their estimates must be interpreted in light of their samples, definitions, and collection methods.

Where observations are limited, structured expert elicitation may be used to establish initial estimates or plausible ranges. Experts should be asked to state the basis of their estimates and the level of confidence they assign to them. Their judgements should be treated as uncertain inputs, not as substitutes for observed data.

3.3 Variables and operational definitions

Table 1 presents the principal variables required by the model.

Variable	Symbol	Operational definition
Threat-event frequency	λ_S	Expected annual frequency of threat events for scenario \(s\), under the selected model
Conditional compromise probability	p_S	Probability that a relevant attempt results in compromise, given the stated conditions
Incident count	N_s	Number of loss-producing incidents for scenario \(s\) during the analysis period
Operational loss	$L_{op,s}$	Financial value of operational disruption attributable to scenario \(s\)
Other financial loss	$L_{ein,s}$	Direct financial consequences not already included in operational or recovery loss
Recovery cost	$C_{rec,s}$	Expenditure required for investigation, containment, restoration, and defined recovery activities
Control effectiveness	Θ_k	Modelled effect of control \(k\) on specified event or loss parameters
Control cost	C_k	Cost of acquiring, implementing, operating, and maintaining control \(k\) over the chosen horizon
Annual aggregate loss	A	Sum of incident losses across the scenarios included in the model
Uncertainty parameters	**khi**	Parameters describing uncertainty in frequency, severity, costs, and dependencies

Table 1. Principal variables in the proposed cyber-risk model.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

The symbols are scenario-specific where appropriate. Definitions must be fixed before data are collected to ensure that different analysts estimate the same quantities consistently.

3.4 Scenario selection

The model should begin with a manageable set of scenarios that are material to the organization. Three illustrative scenarios are proposed.

Scenario A: Ransomware and service interruption. An attacker compromises systems, encrypts or disrupts data, and affects the availability of one or more business services. Losses may include lost operational output, incident response, data restoration, specialist support, and additional continuity costs.

Scenario B: Unauthorized access and data breach. An attacker gains access to sensitive information. Potential consequences include investigation, legal support, customer notification, remediation, regulatory exposure, and other financial consequences supported by the organization's circumstances.

Scenario C: Third-party service outage. A critical external provider experiences an incident that interrupts the organization's operations. The resulting losses may include downtime, alternative-service expenditure, contractual consequences, and additional recovery work.

These scenarios are illustrative rather than a claim that they are the most significant threats for every organization. Actual scenario selection should reflect business-service criticality, threat exposure, regulatory obligations, incident history, and organizational priorities.

3.5 Analysis horizon and parameter estimation

The proposed initial analysis horizon is one year. This period supports the estimation of expected annual loss and comparison of alternative controls using a common timeframe. Multi-year analysis can be added when the organization has sufficient evidence to model changing exposure, investment costs, and control effectiveness.

Incident frequency can be estimated from internal records when those records are sufficiently complete and relevant. For a simple Poisson count model, the estimated rate may be based on observed incidents divided by the number of comparable observation-years. However, the analyst must investigate whether changes in technology, controls, attacker behaviour, and reporting practices make the historical rate unrepresentative of future exposure.

Conditional compromise probability and incident severity may be estimated using relevant data, control testing, threat analysis, or structured expert judgement. Where a point estimate would obscure substantial uncertainty, a probability distribution or plausible interval should be used instead.

The selection of a distribution should be justified by the nature of the variable and the evidence available. It should not be chosen solely because it is convenient for simulation.

3.6 Modelling uncertainty

The model should represent uncertainty in frequency, compromise probability, operational downtime, financial loss, and recovery costs. Probability distributions may be used where the available evidence supports their specification. Where the evidence is too limited to justify a particular distribution, scenario ranges and alternative assumptions should be documented and tested.

A simulation-based approach can repeatedly sample from the chosen input distributions, generate incident counts and severities, and calculate aggregate annual loss. The resulting distribution provides information about the range of possible outcomes under the model's assumptions.

The method should distinguish variability in actual incident outcomes from uncertainty about the parameters and model structure. These may be examined through nested simulation, alternative parameter sets, or separate sensitivity analyses. Saltelli et al. (2008) provide established methods for analysing how variation in model inputs affects output variability.

3.7 Model verification and validation

Verification asks whether the model has been implemented correctly. It should include checks of formulas, units, scenario definitions, loss-category boundaries, and simulated outputs under simple test cases.

Validation asks whether the model adequately represents the real-world process for its intended use. Where suitable historical data exist, predicted incident frequencies and loss distributions can be compared with observed outcomes. Calibration and validation data should be separated where feasible to reduce the risk of judging the model solely on the observations used to construct it.

Additional evaluation should compare the proposed model with a baseline such as a qualitative risk matrix, a simple expected-loss calculation, or an appropriately implemented FAIR-based assessment. The comparison should examine not only numerical estimates but also sensitivity to assumptions, consistency of rankings, interpretability, and usefulness for investment decisions.

3.8 Ethical and practical considerations

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

Incident data may contain sensitive operational details, personal information, or commercially confidential records. Any implementation should apply appropriate access controls, anonymization or aggregation, data-minimization practices, and retention policies.
Model outputs should support, rather than replace, expert judgement and governance. A low estimated annual loss does not necessarily justify accepting a risk when the scenario involves legal obligations, safety implications, critical services, or severe consequences that the model does not capture adequately.

IV. DEVELOPMENT OF THE QUANTITATIVE CYBER-RISK MODEL

4.1 Model structure and assumptions

The proposed model represents cyber risk as the aggregate financial loss generated by a defined set of incident scenarios over a specified period. Let *s = 1, 2… S* denote the incident scenarios, where each scenario represents a particular threat and business-consequence pathway.
Let:

- N_s be the number of loss-producing incidents for scenario \(s\) during the analysis period.
  $L\{_{s,j}\}$ be the financial loss caused by the \(j\)-th incident under scenario \(s\).
- A be total aggregate loss across all scenarios.

Aggregate annual loss is defined as:

$$A = \sum_{\{s=1\}}^{\{S\}} . \sum_{\{j=1\}}^{\{N\}} . Ls, j$$

This formulation allows the model to represent multiple incidents within a year and differences in severity between incidents. It also enables the analysis of losses across multiple scenarios rather than limiting the assessment to a single threat.
The model initially assumes that the analysis period is one year and that the scenarios have sufficiently clear definitions to avoid overlapping incident counts. Where one event produces several consequences, those consequences should be recorded within the same incident-loss calculation unless there is a defensible reason to model them separately.
The model does not require all scenarios to be independent. Dependencies may be represented through shared parameters, conditional probabilities, common-cause events, or joint simulation.

4.2 Modelling threat-event probability

Threat-event frequency and conditional compromise probability must be distinguished. Suppose a threat event occurs with annual frequency λT,s, and the conditional probability that it produces a compromise is p_s. Under a simplified model in which each relevant attempt has a consistent compromise probability and the assumptions of independent event thinning are reasonable, the expected frequency of compromise events is: λ_S C,s λ_S {T,s}p_s
Here, λ_S{C,s} is the expected annual frequency of compromises for scenario *s*.
This relationship is a simplification, not a universal law of cyber incidents. It assumes that the event definition and probability are compatible and that the relevant attempts can be treated consistently. If attempts differ substantially in severity or success probability, or if attackers adapt to controls, a more detailed model is required.
For a scenario that requires several sequential conditions, the probability of compromise can be decomposed into conditional probabilities. For example:

$$[\, P(M\backslash mid\ T) = P(C\backslash mid\ T)\backslash, P(M\backslash mid\ C, T)\ \backslash]$$

where *T* denotes a threat event, *C* denotes an intermediate compromise, and *M* denotes a material business incident. This expression applies to the specified sequence when the conditional events are defined consistently. Other attack pathways may require alternative structures.
Where annual incident counts reasonably follow a Poisson distribution, the probability of observing *n* incidents is:

$$P(N_s = n) = \backslash frac\{e\^{}\{\lambda_s\}\backslash\lambda_s\^{}n\}\{n!\}\ \backslash]$$

In this expression λ_s is the expected annual frequency of the loss-producing incident. The Poisson model assumes a constant rate and independent event occurrences over the period. If incidents cluster, rates change over time, or multiple losses arise from one common event, those assumptions must be reconsidered.

4.3 Integrating operational disruption, financial exposure, and recovery costs

The financial loss associated with each incident is defined as:

$$L_s = L_{\{op,s\}} + L_{\{fin,s\}} + C_{\{rec,s\}}$$

where:

- ($L_{\{op,s\}}$) represents operational disruption losses.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

- ($L_{\{fin,s\}}$) represents other financial consequences.
- ($C_{\{rec,s\}}$) represents recovery expenditure.

Operational loss may be estimated using a service-specific cost rate and the duration of disruption. A simplified representation is: $L_{\{op,s\}} = D_s \times R_s$

where D_s is the effective duration of disruption and R_s is the estimated financial loss per unit of disruption time.

The term effective duration is important because not every minute of technical unavailability produces the same financial consequence. Some services may have workarounds, spare capacity, or the ability to defer transactions without permanently losing them. Consequently, R_s should be estimated from the business process affected rather than selected as a generic rate.

Other financial losses may include legal support, customer compensation, contractual penalties, and other documented costs that do not already appear in operational loss or recovery expenditure. Their inclusion depends on the scenario and the organization's circumstances.

Recovery costs may be represented as:

$$C_{\{rec,s\}} = C_{\{inv,s\}} + C_{\{con,s\}} + C_{\{rest,s\}} + C_{\{val,s\}}$$

where:

- $C_{\{inv,s\}}$ is investigation expenditure.
- $C_{\{con,s\}}$ is containment and eradication expenditure.
- $C_{\{rest,s\}}$ is system and data restoration expenditure.
- $C_{\{val,s\}}$ is post-restoration validation and recovery-related expenditure.

These categories provide a starting structure and may be modified to suit the organization's accounting and incident-response practices. They should be defined so that each expenditure is counted once. For example, if specialist incident-response labour is included in investigation expenditure, it should not also be counted under a separate general recovery-cost category.

The decomposition also distinguishes the economic consequences of disruption from the expenditure required to restore systems. This is particularly important for ransomware incidents, where restoration costs and lost business activity can occur simultaneously.

4.4 Estimating expected annual loss

Expected annual loss (EAL) is the probability-weighted average of aggregate annual loss. Under the basic model, it is expressed as:

$$E[A] = sum\{s = 1\}^\{S\}E[Ns]E[Ls]$$

This expression is valid when the relevant incident-count and incident-severity expectations can be combined as shown—for example, when the expected incident severity is appropriately defined and does not vary with incident count in a way that invalidates the factorization. More generally, the model should retain the conditional expectation:

$$E[A] = sum\{s = 1\}^\{S\}E[Ns_Ls]$$

where L_s represents an appropriately defined incident-level severity measure in the corresponding expectation.

Expected annual loss is useful for comparing scenarios and estimating the average annual financial exposure under the model's assumptions. However, it does not describe the full distribution of possible losses. Two scenarios may have similar expected losses but very different tail behaviour: one may produce frequent, moderate losses, while the other may produce infrequent but catastrophic losses.

For that reason, EAL should be used alongside additional distributional measures when the organization is exposed to severe incidents.

4.5 Loss-exceedance probability and tail risk

The loss-exceedance probability measures the likelihood that aggregate annual loss exceeds a specified threshold \x\: P (A>x)

This measure can help decision-makers examine the probability of losses exceeding an organization's tolerance, insurance threshold, or available financial capacity.

Value at Risk (VaR) can be used to describe a loss quantile at a selected confidence level. Conditional Value at Risk (CVaR), also known in many contexts as expected shortfall, describes the expected loss in the tail beyond the relevant VaR threshold, subject to the definition used.

These measures provide information that EAL alone cannot. Nevertheless, they should be interpreted carefully when severe losses are poorly represented in the available data. A precise-looking tail-risk estimate is not necessarily reliable when its underlying assumptions are weak.

4.6 Representing dependencies between scenarios

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

Cyber incidents can affect multiple systems, business units, and external services simultaneously. A compromise of shared identity infrastructure, for example, may affect several applications. Similarly, an outage at a critical service provider may disrupt multiple business processes.

Assuming complete independence between scenarios could underestimate aggregate loss if common causes produce correlated incidents. Böhme and Schwartz (2010) discuss interdependence and correlated risk in cyber-insurance modelling, while Fedele and Roner (2022) review the importance of interdependent firms in cybersecurity investment research.

The proposed model should represent material dependencies through shared risk factors, conditional event structures, or joint sampling. If a common event affects several systems, the simulation should model the common cause and its consequences rather than independently generating each downstream loss as though the systems had no connection.

Where the data do not support a reliable dependency estimate, the model should test alternative plausible dependency assumptions and report their influence on aggregate loss.

4.7 Monte Carlo simulation

Monte Carlo simulation provides a practical way to propagate uncertainty through the model. In each simulation iteration, the analyst samples relevant parameters, generates incident counts, samples incident severities, calculates operational and financial losses, and aggregates the results.

A simplified implementation proceeds as follows:

1. Define the scenarios, parameter distributions, and dependency structure.
2. Sample incident-frequency and compromise parameters.
3. Generate incident counts for the analysis period.
4. Sample operational disruption, other financial losses, and recovery expenditure for each incident.
5. Calculate incident-level losses and aggregate annual loss.
6. Repeat the process for a sufficiently large number of iterations.
7. Summarize the simulated distribution using EAL, selected quantiles, exceedance probabilities, and uncertainty measures.
8. Repeat the simulation under alternative control and dependency assumptions.

The number of iterations should be chosen based on convergence checks and the precision required for the decision. A large iteration count cannot compensate for poor input data, invalid assumptions, or a misspecified model.

4.8 Modelling control effectiveness and residual risk

Security controls may affect different stages of the incident pathway. Preventive controls can reduce the probability of successful compromise; detection and containment controls may reduce incident severity or duration; and recovery controls may reduce restoration time and expenditure.

Let **x** represent a selected set of controls. The residual loss distribution is then conditional on the controls implemented:

$$A_x \sim P(A \mid x, \varphi)$$

where **φ** represents uncertain model parameters.

The model should not assume that every control reduces all loss components. Instead, each control should be mapped to the parameters it is expected to influence, supported where possible by control-testing results, incident evidence, vendor-independent evaluations, or documented expert assessment.

For instance, a tested backup and restoration process may reduce the distribution of recovery duration and restoration expenditure without necessarily reducing the number of initial attack attempts. Multi-factor authentication may reduce the probability of certain unauthorized-access pathways without directly reducing every possible recovery cost.

This approach allows the model to distinguish preventive effectiveness from resilience and recovery effectiveness.

4.9 Preliminary investment metric

A simple return on security investment (ROSI) measure can be used to screen alternative controls:

$$ROSI_k = (\Delta EAL_k - C_k) / C_k \times 100\%$$

where $\Delta EAL_k$is the estimated reduction in expected annual loss attributable to control $_k$, and C_k is its cost over the same analysis horizon.

This measure is meaningful only when the benefit and cost are expressed on a consistent basis. A one-year expected-loss reduction should not be compared directly with a multi-year implementation cost without appropriate adjustment. The analysis should also avoid treating uncertain loss reduction as guaranteed savings.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

ROSI is therefore a preliminary decision aid, not a complete investment rule. It may omit legal obligations, strategic dependencies, nonfinancial impacts, risk appetite, and extreme-loss considerations. The final decision should use the residual-loss distribution and the organization's wider governance requirements.

V. PROPOSED MODEL IMPLEMENTATION AND EVALUATION

5.1 Implementation workflow

The model can be implemented using a spreadsheet for a simple initial assessment or a statistical programming environment for more complex simulations. The implementation should separate input data, parameter assumptions, loss calculations, control scenarios, and reporting outputs so that the analysis can be reviewed and reproduced.

The following workflow is proposed:

- Stage 1 — Establish the baseline: document existing controls, define scenarios, and estimate current incident frequency and loss distributions.
- Stage 2 — Specify loss categories: establish operational, other financial, and recovery-cost definitions, including rules to prevent double-counting.
- Stage 3 — Represent uncertainty: select defensible parameter distributions or scenario ranges and document the evidence supporting them.
- Stage 4 — Simulate baseline risk: estimate aggregate annual losses, EAL, selected loss quantiles, and exceedance probabilities.
- Stage 5 — Evaluate controls: modify only the parameters expected to change under each control and rerun the simulation.
- Stage 6 — Compare alternatives: assess residual risk, cost, uncertainty, budget requirements, and operational feasibility.
- Stage 7 — Document findings: report assumptions, sensitivity results, limitations, and decision implications.

5.2 Evaluation scenarios

The three scenarios introduced in Section 3 can be used to test whether the framework handles different types of cyber loss consistently.

For ransomware, the analysis would focus on incident frequency, compromise probability, disruption duration, and the cost of investigation and restoration. For a data breach, it would emphasize the number and sensitivity of affected records where relevant, investigation and notification expenditure, legal exposure, and other supported consequences. For a third-party outage, it would examine provider dependency, service interruption, alternative-service costs, and the possibility that one event affects several business processes.

These scenarios are intended to test the model's structure. They do not establish that the framework has been implemented or that one scenario necessarily has a greater expected loss than another.

5.3 Baseline comparisons

The proposed model should be compared against at least two baselines:

1) Qualitative risk matrix: a conventional likelihood-and-impact classification.
2) Simplified expected-loss model: a model that estimates expected loss but omits some combination of detailed recovery costs, uncertainty, or dependencies.

Where data and expertise permit, a FAIR-based assessment can provide an additional comparator. The objective is not to assume that the proposed model is superior, but to determine whether its additional complexity produces decision-relevant benefits.

Evaluation criteria should include the consistency of scenario rankings, sensitivity to uncertain assumptions, ability to represent extreme losses, transparency of cost decomposition, reproducibility, and usefulness to decision-makers.

5.4 Validation criteria

Validation should address both statistical and practical performance.

Statistical evaluation may include comparison between predicted and observed incident frequencies, assessment of predicted loss ranges against observed losses, and testing whether estimated distributions are appropriately calibrated where sufficient data exist.

Decision evaluation may examine whether investment recommendations remain stable under plausible parameter changes, whether the model identifies important dependencies, and whether its outputs change when controls affect only the relevant parts of the incident pathway.

Practical evaluation should assess data requirements, implementation effort, interpretability, and whether decision-makers can understand the consequences of changing assumptions.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

A framework that generates more detailed numbers is not automatically better. Its usefulness depends on whether those numbers are credible, understandable, and relevant to the decisions being made.

VI. RESULTS AND SENSITIVITY ANALYSIS

6.1 Overview of the Results

The proposed cyber-risk quantification model integrates incident probability, operational disruption, direct financial exposure, and recovery expenditure within a unified quantitative framework. The analysis is designed to estimate aggregate annual loss, characterize uncertainty in potential losses, and assess how alternative cybersecurity controls influence residual risk and investment priorities.

The principal outputs comprise expected annual loss (EAL), the distribution of aggregate annual loss, loss-exceedance probabilities, conditional tail-loss measures, and the relative contributions of operational disruption, direct financial exposure, and recovery expenditure. These outputs provide a basis for comparing risk scenarios and evaluating the potential economic benefits of cybersecurity investments.

6.2 Aggregate Annual Loss and Risk Distribution

Aggregate annual loss is defined as the sum of losses associated with all relevant cyber incidents during the analysis period:

$$A = \Sigma(s = 1 \text{ to } S)\ \Sigma(j = 1 \text{ to } N_s)\ L_{s,j}$$

where A represents aggregate annual loss, S denotes the number of cyber-risk scenarios, N_s represents the number of incidents associated with scenario s, and $L_{s,j}$ denotes the loss arising from incident j under scenario s.

The model estimates expected annual loss alongside the median and selected upper quantiles of the aggregate-loss distribution. Expected annual loss provides a measure of average financial exposure, while upper quantiles and loss-exceedance probabilities characterize the potential for less frequent but more severe outcomes.

Loss-exceedance analysis evaluates the probability that aggregate annual loss exceeds a specified financial threshold. Conditional tail-loss measures can further characterize the expected severity of losses within the upper tail of the distribution. Together, these measures support risk prioritization beyond reliance on average annual loss alone.

6.3 Contributions of Operational Disruption, Financial Exposure, and Recovery Costs

Incident-level loss is decomposed into three principal components:

$$L_s = L_{op,s} + Lfin_{,s} + C_{rec,s}$$

where $L_{op,s}$ represents operational disruption losses, $Lfin_{,s}$ denotes other direct financial exposure, and $C_{rec,s}$ represents recovery expenditure associated with scenario s.

This decomposition enables the model to distinguish losses arising from interrupted business operations from direct financial exposure and the expenditure required to restore affected systems and services. It also provides a basis for assessing whether recovery costs materially influence total estimated losses or alter the relative importance of individual cyber-risk scenarios.

The contribution of each component should be assessed in absolute monetary terms and, where appropriate, as a proportion of total estimated loss. Such comparisons can help decision-makers identify whether their principal exposure arises from operational downtime, direct financial consequences, recovery requirements, or a combination of these factors.

6.4 Sensitivity Analysis and Parameter Uncertainty

Sensitivity analysis examines the extent to which changes in uncertain model inputs influence estimated losses and cybersecurity investment priorities. The principal parameters include incident frequency, compromise probability, operational disruption duration, loss per unit of disruption, recovery expenditure, and control effectiveness.

One-at-a-time sensitivity analysis evaluates the effect of varying individual parameters while holding other inputs constant. Although this approach provides an accessible assessment of parameter influence, it may not adequately capture interactions between uncertain variables. Global sensitivity analysis, following the principles discussed by Saltelli et al. (2008), can provide a broader assessment of the combined effects of multiple uncertain inputs.

Particular attention should be given to interactions between disruption duration and business-service criticality, as well as between compromise probability and the severity of potential consequences. Shared infrastructure, common security controls, and other dependencies should also be considered because they may affect aggregate losses and the relative performance of alternative controls.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

The sensitivity results should identify the parameters associated with the greatest changes in expected annual loss, upper-tail risk measures, and control rankings. This information can guide further data collection and help determine which assumptions require more precise estimation.

6.5 Implications for Cybersecurity Investment Decisions

The model connects quantified residual risk with investment prioritization by comparing estimated losses before and after the introduction of alternative controls. The assessment should consider both the reduction in expected annual loss and the change in the distribution of potential losses.

Return on security investment (ROSI) can be expressed as:

$$ROSI_k = [(\Delta EAL_k - C_k) / C_k] \times 100\%$$

where ΔEAL_k represents the estimated reduction in expected annual loss attributable to control k, and C_k denotes the cost of implementing and maintaining that control over a consistent evaluation period.

Investment comparisons should account for implementation costs, recurring expenditure, uncertainty in control effectiveness, and the possibility that a control reduces some risk scenarios more effectively than others. Where control rankings vary substantially across plausible assumptions, decision-makers should consider the robustness of each option rather than relying exclusively on a single point estimate.

6.6 Summary of the Results Framework

The integrated framework provides a structured basis for evaluating cyber risk through incident probability, financial consequences, recovery expenditure, and uncertainty. Its principal analytical contribution is the alignment of risk estimation with investment decisions, enabling organizations to assess both expected losses and the potential consequences of severe events.

The practical value of the framework depends on the quality of its input data, the suitability of the selected probability distributions, the treatment of dependencies, and the validation of model assumptions. Numerical conclusions about risk reduction, sensitivity rankings, or the superiority of particular cybersecurity investments should therefore be established from the completed simulation and reported with their corresponding assumptions and uncertainty measures.

VII. CYBERSECURITY INVESTMENT AND RISK PRIORITISATION

7.1 Prioritising controls by their effect on risk

Cybersecurity investment decisions should consider the expected reduction in loss, the distribution of residual risk, the cost of implementation, and the organization's obligations and operational priorities. The framework can help decision-makers distinguish between controls that prevent incidents and controls that reduce their consequences after an incident occurs.

Preventive controls include identity and access management, multi-factor authentication, secure configuration, patch management, and vulnerability remediation. Detection and containment controls include monitoring, endpoint protection, incident-response capabilities, and processes for isolating affected systems. Recovery-oriented controls include tested backups, restoration procedures, contingency arrangements, and alternative service capabilities.

The relevant question is not simply whether a control is effective in general. It is how the control changes the parameters of a particular scenario, how much uncertainty remains, and whether the resulting risk reduction justifies its cost.

Gordon and Loeb (2002) provide a foundational economic perspective on security investment, while Fedele and Roner (2022) demonstrate the breadth of research into firms' incentives to invest in cybersecurity. The proposed model extends this decision logic by connecting control effects to multiple loss categories and to the distribution of residual losses.

7.2 Budget-constrained investment selection

Organizations rarely have unlimited resources to address every identified risk. The model can represent control selection as a budget-constrained optimization problem.

Let x_k represent whether control k is selected, with $x_k = 1$ when selected and $x_k = 0$ otherwise. Let C_k be its cost and B the available budget. The basic constraint is:

$$\sum_{k=1}^{K} C_k x_k \leq B$$

A simple objective is to select the combination of controls that minimizes expected residual loss:

$$\min_x E[A_x]$$

subject to the budget constraint and any required operational or policy constraints.

This formulation provides a starting point, but minimizing expected loss may not be appropriate in every situation. An organization may also need to limit the probability of a loss exceeding a specified threshold, satisfy mandatory controls, preserve service availability, or address risks that cannot reasonably be accepted.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

A more comprehensive decision process can therefore include constraints on tail risk, control dependencies, minimum security requirements, and implementation feasibility. Control costs should include relevant implementation and ongoing operating expenses over the chosen analysis horizon.

7.3 Incorporating risk appetite and nonfinancial constraints

Risk appetite determines the amount and type of residual risk an organization is prepared to accept. The proposed model can support this decision by estimating the likelihood and consequences of different loss outcomes, but the model cannot determine acceptable risk independently of organizational objectives and obligations.

For example, two scenarios with similar EAL may differ substantially in their potential effects on safety, legal compliance, sensitive information, or essential services. These consequences may require additional constraints or separate assessment rather than being reduced to a single monetary value.

7.4 The value of improved information

Sometimes the most useful next step is not immediate investment in another control but reducing uncertainty about a material risk. An organization might improve incident logging, conduct restoration tests, obtain better service-dependency information, or collect more accurate estimates of business interruption costs.

The value of such information depends on whether it can materially change a decision. If a control remains the preferred option across all plausible parameter values, additional data may have limited immediate decision value. If the ranking changes significantly with uncertain assumptions, further investigation may be worthwhile.

This consideration is particularly important where cyber-risk data are incomplete, as discussed by Cremer et al. (2022). It also helps avoid a false choice between acting immediately and waiting for perfect data: organizations can implement necessary controls while improving the evidence used for subsequent decisions.

VIII. DISCUSSION

8.1 Integrating technical risk with business consequences

The central argument of this paper is that cyber-risk assessment should connect threat events to the business consequences they can produce. A technical vulnerability does not translate directly into a fixed financial loss. Its consequences depend on whether it is exploited, the systems affected, the controls in place, the duration of disruption, and the organization's ability to respond and recover.

NIST (2012) and The Open Group (2020) provide established foundations for risk assessment and quantitative information-risk analysis. The present framework combines these principles with explicit business-loss categories, uncertainty representation, and a structured investment decision process.

The framework should not be interpreted as replacing established standards. Instead, it offers a model-development structure for organizations seeking to connect risk analysis more directly with financial and operational decisions. Its practical value will depend on the quality of its implementation and the suitability of its assumptions.

8.2 Why recovery costs deserve explicit treatment

Recovery expenditure can be significant, but it may be difficult to estimate consistently because organizations record response and restoration costs differently. Some expenses are visible in incident-response invoices, while others appear as staff time, infrastructure expenditure, replacement services, or delayed projects.

Incident-handling and contingency-planning guidance from Cichonski et al. (2012) and Swanson et al. (2010) provides a basis for identifying the activities that should be considered. The model then translates these activities into financial categories, subject to clear boundaries and supporting evidence.

Explicitly modelling recovery costs can improve transparency by showing decision-makers how much estimated loss arises from service disruption, direct financial consequences, and restoration activity. However, it does not guarantee a higher or more accurate estimate in every case. The effect depends on whether recovery costs were previously omitted, included elsewhere, or measured incorrectly.

8.3 Uncertainty as a decision variable

A single expected-loss estimate can conceal substantial variation in possible outcomes. An organization facing a relatively predictable pattern of moderate losses may require a different response from one exposed to a small probability of a severe disruption, even if their EAL estimates are similar.

The framework therefore combines expected annual loss with loss-exceedance probabilities, selected quantiles, and sensitivity analysis. These measures help decision-makers understand both the average exposure and the potential for outcomes that exceed financial or operational tolerance.

Nevertheless, probability distributions are only as defensible as their supporting evidence. Where incident data are scarce or definitions differ, the model should communicate uncertainty explicitly. A more complex model may be less useful than a simpler one if its parameters cannot be supported or its outputs cannot be interpreted.

8.4 Implications for cybersecurity investment

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

A quantitative framework can make the logic of investment decisions more transparent. It can show whether a proposed control is expected to reduce incident frequency, reduce the severity of successful incidents, shorten recovery time, or limit a particular loss category.

This distinction matters because different controls may address different parts of the incident pathway. A preventive control may reduce compromise probability without significantly changing recovery costs conditional on an incident. A recovery control may have limited influence on incident frequency while substantially reducing the consequences of an outage.

The economic literature on cybersecurity investment, including Gordon and Loeb (2002) and Fedele and Roner (2022), provides a basis for analysing such trade-offs. However, the framework should not equate estimated monetary benefits with the entirety of security value. Compliance requirements, strategic objectives, nonfinancial harm, and operational dependencies may impose additional constraints.

8.5 The role of third-party dependencies

Organizations increasingly depend on external providers for cloud hosting, identity management, payment processing, communications, and other critical services. A provider incident can affect several business functions at once, making the organization's exposure dependent on risks beyond its direct control.

Böhme and Schwartz (2010) discuss correlated risks and interdependence in cyber-insurance modelling, while Fedele and Roner (2022) review cybersecurity investment models involving interdependent firms. These perspectives support the inclusion of shared dependencies in the proposed framework.

In practice, dependency modelling requires an inventory of critical providers and services, an understanding of the business processes they support, and an assessment of alternative arrangements. If these relationships are ignored, the model may understate the possibility that a single event produces losses across several scenarios.

8.6 Theoretical and practical contribution

The proposed framework makes three principal contributions.

First, it provides a common structure for representing threat-event frequency, conditional compromise, operational disruption, other financial losses, and recovery expenditure. This makes the composition of risk estimates easier to examine.

Second, it places uncertainty and dependencies within the model rather than treating them as footnotes to a single numerical estimate. This encourages the reporting of ranges, distributions, and sensitivity results.

Third, it connects risk estimation with investment selection under budget constraints. Rather than stopping at the ranking of threats, the framework asks how different controls change residual loss and how those changes compare with implementation and operating costs.

These are conceptual contributions. Their practical significance must be established through implementation, comparison with existing approaches, and empirical validation.

8.7 Limitations

Several limitations should be acknowledged.

Data quality and availability. Incident records may be incomplete, inconsistent, or biased toward events that are detected and reported. External studies may use different definitions of incidents and losses.

Difficulty of financial valuation. Reputation, customer confidence, long-term strategic harm, and some regulatory consequences may be difficult to represent accurately as monetary values. These effects should not be assigned arbitrary values merely to fit the model.

Model assumptions. Probability distributions, incident independence, and dependency structures may not reflect the actual threat environment. The assumptions must be tested and revised when evidence warrants.

Changing threat conditions. Attack methods, technology, controls, and organizational exposure change over time. Parameters estimated from historical observations may become outdated.

Implementation complexity. A model with many variables can become difficult to maintain and explain. Its complexity should be proportionate to the importance of the decision, the availability of evidence, and the organization's analytical capacity.

Incomplete representation of nonfinancial constraints. Some risks cannot be managed through monetary optimization alone. Legal requirements, safety, ethical obligations, and essential-service responsibilities may require explicit controls regardless of the model's EAL estimate.

IX. CONCLUSION AND RECOMMENDATIONS

9.1 Conclusion

This paper has developed a conceptual quantitative framework for integrating threat probability, business impact, and recovery costs under uncertainty. The framework represents cyber risk through incident frequency and

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

severity, decomposes incident losses into operational disruption, other financial consequences, and recovery expenditure, and aggregates those losses across defined scenarios.

Expected annual loss provides a baseline measure for comparing scenarios, while loss-exceedance probabilities and tail-risk measures help describe outcomes that may be concealed by averages. Monte Carlo simulation offers a proposed method for propagating uncertainty through the model, and explicit dependency structures can account for incidents that affect several systems or services simultaneously.

The framework also links residual-risk estimates to cybersecurity investment decisions. By mapping controls to the parameters they influence and comparing estimated loss reduction with implementation and operating costs, organizations can make investment assumptions more transparent.

However, the framework remains a proposed model rather than an empirically validated solution. Its effectiveness cannot be established without suitable data, implementation, sensitivity analysis, and comparison with existing methods. The contribution of this paper is therefore a structured basis for further research and practical testing, not a claim that a particular level of predictive or financial improvement has already been achieved.

9.2 Recommendations

Based on the proposed framework, the following recommendations are made.

1. Standardize incident and loss recording. Organizations should maintain consistent definitions for incident frequency, operational disruption, direct financial exposure, and recovery expenditure.
2. Separate loss categories explicitly. Analysts should document which costs are included in each category and prevent double-counting across operational losses, response costs, and restoration expenditure.
3. Represent uncertainty transparently. Where data are incomplete, estimates should be accompanied by ranges, distributions, evidence sources, and explanations of key assumptions.
4. Evaluate both average and severe losses. Expected annual loss should be supplemented by exceedance probabilities or tail-risk measures when the organization is exposed to material extreme outcomes.
5. Map controls to specific risk mechanisms. Preventive, detection, containment, and recovery controls should be assessed according to the parameters they are expected to influence.
6. Include external dependencies. Risk assessments should consider critical providers, shared infrastructure, and common-cause events that may affect several business functions.
7. Validate the model before relying on its rankings. Organizations should verify implementation, compare estimates with observed outcomes where possible, and test whether investment recommendations remain stable under plausible assumptions.
8. Integrate quantitative outputs with governance. Investment decisions should also account for compliance, risk appetite, operational resilience, and nonfinancial consequences that cannot be represented adequately by monetary estimates alone.

9.3 Directions for future research

Future research should implement the framework using real organizational incident and financial data. Studies could compare its performance with qualitative risk matrices, simplified expected-loss models, and established quantitative approaches such as FAIR.

Further work could investigate Bayesian estimation where incident observations are sparse, alternative frequency distributions where incidents cluster, and more sophisticated dependency structures for shared providers and infrastructure. Research could also evaluate how restoration testing and operational resilience investments change recovery-time distributions and total losses.

Finally, longitudinal studies could assess whether risk estimates remain calibrated as organizations introduce new controls and as threat conditions evolve. Such research would help determine whether the framework improves risk prioritization, produces more robust investment decisions, or simply adds complexity without sufficient practical benefit.

REFERENCES

1) Anderson, R. (2001). Why information security is hard—An economic perspective. In Proceedings of the 17th Annual Computer Security Applications Conference (pp. 358–365). IEEE. https://doi.org/10.1109/ACSAC.2001.991552
2) Anderson, R., & Moore, T. (2006). The economics of information security. Science, 314(5799), 610–613. https://doi.org/10.1126/science.1130992
3) Böhme, R., & Kataria, G. (2006). Models and measures for correlation in cyber-insurance. In Proceedings of the Fifth Workshop on the Economics of Information Security (WEIS 2006).
4) Böhme, R., & Schwartz, G. (2010). Modeling cyber-insurance: Towards a unifying framework. Proceedings of the Ninth Workshop on the Economics of Information Security (WEIS 2010).

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

5) Cavusoglu, H., Mishra, B., & Raghunathan, S. (2004). The effect of Internet security breach announcements on market value: Capital market reactions for breached firms and Internet security developers. International Journal of Electronic Commerce, 9(1), 70–104. https://doi.org/10.1080/10864415.2004.11044320
6) Cichonski, P., Millar, T., Grance, T., & Scarfone, K. (2012). Computer security incident handling guide (NIST Special Publication 800-61 Rev. 2). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-61r2
7) Cremer, F., Sheehan, B., Fortmann, M., Kia, A. N., Mullins, M., Murphy, F., & Materne, S. (2022). Cyber risk and cybersecurity: A systematic review of data availability. The Geneva Papers on Risk and Insurance—Issues and Practice, 47, 698–736. https://doi.org/10.1057/s41288-022-00266-6
8) European Union Agency for Cybersecurity. (2023). ENISA threat landscape 2023. ENISA.
9) Fedele, A., & Roner, C. (2022). Dangerous games: A literature review on cybersecurity investments. Journal of Economic Surveys, 36(1), 157–187. https://doi.org/10.1111/joes.12456
10) Freund, J., & Jones, J. (2014). Measuring and managing information risk: A FAIR approach. Butterworth-Heinemann.
11) Gordon, L. A., & Loeb, M. P. (2002). The economics of information security investment. ACM Transactions on Information and System Security, 5(4), 438–457. https://doi.org/10.1145/581271.581274
12) Hubbard, D. W., & Seiersen, R. (2016). How to measure anything in cybersecurity risk. Wiley.
13) International Organization for Standardization. (2018). ISO 31000:2018: Risk management—Guidelines. ISO.
14) International Organization for Standardization, & International Electrotechnical Commission. (2022). ISO/IEC 27005:2022: Information security, cybersecurity and privacy protection—Guidance on managing information security risks. ISO.
15) IBM Security. (2023). Cost of a data breach report 2023. IBM.
16) Joint Committee for Guides in Metrology. (2008). Evaluation of measurement data—Guide to the expression of uncertainty in measurement (JCGM 100:2008).
17) National Institute of Standards and Technology. (2012). Guide for conducting risk assessments (NIST Special Publication 800-30 Rev. 1). U.S. Department of Commerce. https://doi.org/10.6028/NIST.SP.800-30r1
18) National Institute of Standards and Technology. (2018a). Framework for improving critical infrastructure cybersecurity (Version 1.1). U.S. Department of Commerce.
19) National Institute of Standards and Technology. (2018b). Risk management framework for information systems and organizations: A system life cycle approach for security and privacy (NIST Special Publication 800-37 Rev. 2). U.S. Department of Commerce. https://doi.org/10.6028/NIST.SP.800-37r2
20) National Institute of Standards and Technology. (2020). Security and privacy controls for information systems and organizations (NIST Special Publication 800-53 Rev. 5). U.S. Department of Commerce. https://doi.org/10.6028/NIST.SP.800-53r5
21) Romanosky, S. (2016). Examining the costs and causes of cyber incidents. Journal of Cybersecurity, 2(2), 121–135. https://doi.org/10.1093/cybsec/tyw001
22) Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., & Tarantola, S. (2008). Global sensitivity analysis: The primer. Wiley.
23) Shostack, A. (2014). Threat modeling: Designing for security. Wiley.
24) Swanson, M., Bowen, P., Phillips, A. W., Gallup, D., & Lynes, D. (2010). Contingency planning guide for federal information systems (NIST Special Publication 800-34 Rev. 1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-34r1

 © Author(s) | CC BY 4.0 International License.