International Journal of
Science, Architecture, Technology and Environment

A Financial-Based Model for Quantifying Cybersecurity Risk Exposure in Enterprise and Digital Trade Infrastructure

The Financial-Based Cybersecurity Risk Exposure Model (FBCREM)

Ayomipo Hannah Alademehin[1*]

[1]*Lamar University, Beaumont, Texas, USA*

[*]*Corresponding author*

DOI: https://doi.org/10.63680/ijsate062683.95

Abstract

Cybersecurity risk has evolved from a discrete technical concern into a material financial liability with computable, auditable consequences for enterprise valuation, regulatory capital adequacy, and the continuity of digital trade infrastructure. Despite this evolution, the dominant risk quantification paradigms, including the National Institute of Standards and Technology Cybersecurity Framework, the International Organization for Standardization guidance on information security risk management, the Control Objectives for Information and Related Technologies governance framework, and the Factor Analysis of Information Risk model, produce qualitative or semi-quantitative outputs that cannot be integrated directly into financial statements, regulatory disclosures, or actuarially grounded insurance pricing. This paper develops the Financial-Based Cybersecurity Risk Exposure Model, a quantitative framework that integrates four analytical components: a probabilistic Annualized Loss Expectancy derived from Monte Carlo simulation across ten thousand iterations; Sector Exposure Multipliers calibrated against the critical infrastructure classifications maintained by the Cybersecurity and Infrastructure Security Agency; Digital Trade Disruption Coefficients that model supply chain cyber loss propagation; and a Regulatory Cost Component derived from six hundred and eighty enforcement actions across five United States regulatory jurisdictions. The combined output is a Cyber Risk Exposure Score expressed in United States dollars at specified confidence levels. Validation against one thousand two hundred and forty-seven enterprise breach incidents across three United States enterprise archetypes demonstrates a mean absolute error of 8.3 percent, outperforming benchmarks based on the Factor Analysis of Information Risk model by twenty-two percentage points and operational risk value-at-risk models by eleven points. This expanded edition adds fully worked numerical examples for each archetype, a formally specified extension for nation-state threats, an implementable specification of a reduced-data variant for small and medium enterprises, an enlarged treatment of artificial intelligence as both an offensive and defensive force in financial cybersecurity, and a deepened analysis of the United States regulatory landscape. The framework addresses a documented gap in enterprise risk management, cyber insurance underwriting, regulatory capital allocation, and digital trade policy, providing a replicable methodology grounded in United States regulatory and market data.

Keywords: cybersecurity risk quantification; Cyber Risk Exposure Score; financial loss expectancy; Factor Analysis of Information Risk; critical infrastructure; Monte Carlo simulation; digital trade security; enterprise risk management; cyber insurance; Annualized Loss Expectancy; securities disclosure; bank capital; operational resilience; artificial intelligence governance

1. Introduction

The global cost of cybercrime has been projected to reach figures measured in the tens of trillions of United States dollars each year, a transfer of economic value without precedent in recorded history and one that, on the most widely cited estimates, exceeds the combined annual output of several of the world's largest economies (Morgan, 2020; Cybersecurity Ventures, 2024). That figure is not a single number but a composite of many distinct harms. It includes the direct financial losses that follow data breaches, the extortion payments demanded by ransomware operators, the funds drained through business email compromise, and the value lost when intellectual property is stolen. It also includes a wide range of indirect costs that are harder to see but no less real: the slow erosion of a brand after a public breach, the regulatory sanctions that follow a compliance failure, the higher insurance premiums that a damaged risk profile attracts, the cost of rebuilding compromised technology, and the quiet withdrawal of customer and investor confidence that can outlast every technical remedy. For the boards, chief financial officers, and risk committees of enterprises that operate within a United States regulatory environment now explicitly demanding financially quantified disclosure of cyber risk, the journey of cybersecurity from a technical inconvenience to a material financial liability is complete.

The question that confronts these decision-makers is therefore no longer whether cybersecurity constitutes a financial risk. That argument has been settled by events. The question is whether organizations possess analytical tools capable of measuring that risk with the precision that investment-grade decision-making demands. It is here that a striking gap appears. The instruments most widely used to manage cyber risk were not designed to produce financial measurements at all, and the consequences of that mismatch are now visible across boardrooms, insurance markets, regulatory filings, and the courts.

Figure 1. Global annual cybercrime cost trajectory, 2015 to 2025. Sources: Cybercrime Magazine (2024); World Economic Forum Global Cybersecurity Outlook (2024). Projected values for the final years are based on the compound annual growth rate of the prior four-year period. [Figure to be inserted from the original manuscript.]

Consider the predominant analytical tools available to the enterprise risk practitioner. The Cybersecurity Framework published by the National Institute of Standards and Technology, now in its second major version and adopted globally as a guide to operational maturity, explicitly declines to provide a method for financial measurement; it confines itself to function-based maturity ratings that describe how well an organization performs a set of activities. The information security management standards published by the International Organization for Standardization, together with their companion guidance on managing information security risk, provide a systematic lifecycle for treating risk, but they retain ordinal rankings, the familiar labels of high, medium, and low, that cannot be entered into a balance sheet or fed into a regulatory capital calculation. The Control Objectives for Information and Related Technologies framework operates at the level of governance capability and process maturity rather than financial quantification. Even the Factor Analysis of Information Risk model, the most quantitatively advanced of the mainstream approaches, relies primarily on the elicitation of probability estimates from subject-matter experts, and it

does not incorporate the sector-specific implications for regulatory capital or the propagation of losses through digital trade networks that characterize the most financially consequential modern cyber incidents.

This misalignment between financial materiality and measurement adequacy is not an abstract concern. It carries real and recurring consequences. When the breach of a major health-care payments processor in 2024 generated losses exceeding one and a half billion dollars and disrupted payment processing for hundreds of thousands of providers, no existing risk quantification framework had predicted a loss of that magnitude for the parent organization's cyber exposure profile. When the United States Securities and Exchange Commission brought its first cybersecurity disclosure enforcement action against a major software company in late 2023, that company had been providing precisely the kind of qualitative, non-financial risk disclosure that a very large share of public companies still produces today. When destructive malware generated roughly ten billion dollars in global losses in 2017, insurers who had priced cyber policies using questionnaire-based actuarial models faced unexpected catastrophic claims that triggered widespread coverage disputes, several of which were resolved only after years of litigation. The pattern is consistent and instructive: qualitative cyber risk frameworks systematically underestimate the financial magnitude of cyber exposure, and the gap between estimated and realized losses carries direct consequences for boards, insurers, regulators, and investors alike.

1.1 The Cost of Imprecision

It is worth dwelling on why imprecision is so costly in this domain, because the answer motivates the entire design of the model developed in this paper. In most areas of enterprise risk, the financial consequences of a hazard are expressed in money as a matter of course. Credit risk is measured in expected and unexpected losses, which are denominated in currency. Market risk is measured through value- at-risk and expected shortfall, again in currency. Liquidity risk is expressed in funding gaps and survival horizons measured in days and dollars. Against this backdrop, cyber risk has been an anomaly, the one major category of enterprise risk routinely expressed in colors and adjectives rather than in money. A board that can compare its credit value-at-risk against its market value-at-risk cannot place its cyber risk on the same axis when that risk arrives as a heat map shaded red, amber, and green.

The cost of this anomaly compounds in three ways. First, it distorts the allocation of capital and attention. When cyber risk cannot be compared in financial terms against other risks, security investment is set by benchmarking against peers, by the persuasive force of the most recent incident in the news, or by the intuition of whichever executive holds the budget, rather than by a disciplined comparison of marginal risk reduction against marginal cost. Second, it weakens the quality of transfer. Insurance exists to move risk from those who bear it to those better placed to pool it, but a market that cannot price the risk it transfers will oscillate between underpricing that invites catastrophic loss and overpricing that drives sound buyers away, and both failures have been visible in the cyber insurance market in recent years. Third, it undermines accountability. Regulators, auditors, and investors increasingly expect that an enterprise can demonstrate it has measured and understood its material risks, and a qualitative rating cannot bear the evidentiary weight that a financially denominated estimate, with stated assumptions and confidence intervals, can carry.

1.2 Research Motivation and Contribution

The Financial-Based Cybersecurity Risk Exposure Model is motivated by three converging developments in the United States regulatory and market landscape, each of which sharpens the demand for financially

denominated cyber risk estimates. The first is the set of cybersecurity disclosure rules adopted by the Securities and Exchange Commission, effective in December 2023, which require registered entities to disclose material cybersecurity incidents within four business days and to describe their risk management processes in annual filings. Materiality is, at its root, a financial concept. It asks whether a reasonable investor would consider an incident important to an investment decision and answering that question responsibly requires risk estimates expressed in money, which the prevailing frameworks cannot provide. The second development is the expansion of cyber stress-testing expectations for large financial institutions, under which supervisors increasingly demand quantitative estimates of losses under specified hypothetical

cyber scenarios, estimates that qualitative assessments cannot generate. The third is a line of insurance litigation that has established that the exclusion clauses insurers rely upon to deny certain cyber claims require actuarially sound loss modelling capable of distinguishing among classes of incident, a capability that questionnaire-based underwriting plainly lacks.

Against this backdrop, the model developed here makes four specific contributions. It offers the first formally specified financial cyber risk quantification model to integrate sector-calibrated multipliers, supply-chain propagation coefficients, and regulatory cost components within a single auditable equation. It grounds its validation in realized United States enterprise losses rather than in expert elicitation. It treats the disruption of digital trade as a first-class analytical component rather than an afterthought. And it is designed for direct applicability to the contexts of securities disclosure, bank capital adequacy, and operational resilience compliance. This expanded edition adds to those contributions in five further directions, each developed at length in the sections that follow: fully worked numerical examples that trace the calculation from inputs to final score for each enterprise archetype; a formally specified extension that brings nation-state threats into the model as a distinct term; an implementable specification of a reduced-data variant suitable for small and medium enterprises that lack access to commercial data subscriptions; an enlarged treatment of artificial intelligence as a force that reshapes both the offensive threat and the defensive response in financial cybersecurity; and a deepened analysis of how the model maps onto the specific obligations created by United States and international regulation.

A note on scope and intent

This paper develops a methodology and validates it against historical data. It does not offer investment, legal, or actuarial advice, and the figures it reports are drawn from the sources cited rather than from any single proprietary dataset belonging to one institution. Practitioners who adopt the model should calibrate it against their own data and seek qualified professional review before relying on its outputs for regulatory, capital, or insurance decisions.

1.3Structure of the Paper

The remainder of the paper proceeds as follows. Section 2 reviews literature on cybersecurity risk quantification and establishes the theoretical foundations on which the model rests, tracing three generations of thought and identifying five structural limitations of existing frameworks. Section 3 specifies the mathematical form of the model in full, component by component. Section 4 presents fully worked numerical examples that trace the calculation for each of the three enterprise archetypes, so that the reader can follow the arithmetic from raw inputs to the final score. Section 5 validates the model against realized

United States losses and benchmarks its accuracy against established alternatives. Section 6 develops the extension of the model to nation-state threats, converting what earlier treatments left as a future-work note into a formally specified additional term. Section 7 conducts a global sensitivity analysis and sets out the model's limitations candidly. Section 8 develops the practical applications of the model across enterprise risk management, insurance, regulatory capital, and digital trade policy. Section 9 provides detailed implementation guidance, including a maturity model and effort estimates. Section 10 specifies a reduced-data variant for smaller organizations. Section 11 examines the role of artificial intelligence on both sides of the contest between attackers and defenders. Section 12 sets out directions for future research, and Section 13 concludes.

1.4Who the Model Serves

It helps to be explicit about the constituencies the model is meant to serve, because each brings a different question to it and the model is designed to answer all of them from a single parameterization. The board and its risk committee raise the question of whether the enterprise is carrying an acceptable level of cyber risk relative to its appetite and its capital, and they need an answer in the same financial terms in which they consider every other risk. The chief financial officer raises the question of how much to provision and how to weigh security investment against competing demands and needs marginal financial returns rather than maturity ratings. The chief information security officer brings the question of which controls reduce risk most per dollar and needs defensible prioritization. The insurer raises the question of how to price and limit coverage and needs an entity-specific distribution rather than a sector average. The supervisor raises the question of whether the enterprise has measured and provisioned against its material risks, and needs a reproducible, auditable estimate. And the policymaker brings the question of where systemic risk concentrates in the national digital infrastructure and needs an aggregate the model can construct. A single model that speaks to all these constituencies in their common language, money, is more valuable than a collection of specialized tools that each serve one and cannot be reconciled with the others.

1.5What the Model Does Not Claim

Intellectual honesty requires being as clear about the model's limits as about its capabilities and stating those limits at the outset rather than burying them. The model does not claim to predict whether a particular enterprise will suffer a particular incident in a particular year; it estimates a distribution of possible annual losses, not a forecast of events. It does not claim that its parameters are immutable; they are estimated from historical data and must be recalibrated as conditions change. It does not claim to capture every conceivable source of loss; novel attack types and unprecedented events may fall outside the historical distributions on which it is calibrated, which is why the heavy-tailed distributional choice and the separate treatment of nation-state events are deliberate hedges against the unseen. And it does not claim to replace human judgment; it is an instrument that informs decisions made by accountable people, not an oracle that makes those decisions. A model offered with these limits stated plainly is more trustworthy, not less, than one offered with claims of certainty it cannot support.

1.6The Anatomy of a Cyber Loss

It is worth pausing, before the technical development begins, to trace how a single cyber incident generates financial harm, because the anatomy of a loss is the anatomy of the model. When an enterprise suffers a

breach, the harm arrives in waves. The first wave is the direct cost: the expense of investigating the incident, restoring systems, notifying affected parties, and replacing compromised technology. This is the harm that the direct-loss term of the model captures, and it is the harm that existing frameworks, when they quantify anything, tend to quantify. The second wave is the sector-specific consequence: the breach of a financial institution triggers regulatory and systemic effects that the breach of a less consequential enterprise does not, and the sector multiplier captures this differential. The third wave is propagation: if the breached enterprise sits within a digital trade network, the disruption spreads to partners, suppliers, and customers, generating harm far from the point of compromise, and the trade disruption term captures this wave. The fourth wave is the regulatory response: investigations open, penalties are assessed, remediation is mandated, and, for a bank, capital requirements rise, and the regulatory cost component captures this final wave.

The waves do not arrive at once, and they do not subside at the same rate. The direct cost is felt immediately; the propagation unfolds over days; the regulatory response can take years. A model that captured only the first wave would measure only the visible tip of the harm, and it is precisely because existing frameworks tend to measure only that tip that they systematically understate exposure. The four-component structure of the model is, at bottom, an attempt to measure all four waves, and the validation finding that the later waves often dominate the total is the empirical vindication of that structure. To understand the model is to understand that a cyber loss is not a single event with a single cost, but a cascade whose later stages frequently exceed its beginning.

2. Literature Review and Theoretical Foundations

The intellectual history of cybersecurity risk quantification can be read as a steady migration from the language of compliance toward the language of finance. Each generation of thought has responded to the threats and the regulatory pressures of its era, and each has carried forward both the strengths and the blind spots of its predecessors. Understanding this lineage is essential, because the model developed in this paper is not a rejection of what came before but an attempt to complete an unfinished trajectory, carrying cyber risk measurement the final distance into the financial domain where modern decision-making occurs.

2.1Three Generations of Cyber Risk Quantification

The first generation, spanning roughly the years from 1990 to 2005, was characterized by qualitative, compliance-driven frameworks that emerged from the discipline of information security management. The British standard, later codified internationally as the information security management family, introduced systematic processes for managing information security, while the early federal risk assessment guidance published by the National Institute of Standards and Technology provided one of the first structured methodologies for assessing risk in government systems. These frameworks expressed risk as a qualitative combination of the likelihood of a threat and the criticality of an asset, without translating either into financial loss. They were built to demonstrate compliance, not to inform investment-grade decisions, and they reflected an era in which cyber incidents were understood primarily as operational and reputational events rather than as material financial ones.

The second generation, from roughly 2005 to 2015, introduced probabilistic reasoning. The standardized scoring of technical vulnerability severity provided a common vocabulary for describing how dangerous a given weakness might be, though it offered no path to financial translation. The most significant development

of this period was the Factor Analysis of Information Risk model, developed within the insurance industry and later formalized through a standards body. This approach decomposed cyber risk into the frequency of loss events and the magnitude of loss, broke each of those factors into calibratable sub-components, and advocated the use of Monte Carlo simulation to build distributions over expected losses (Jones and Freund, 2015). It represented a genuine paradigm shift, the first widely adopted attempt to treat cyber risk as a quantity to be modelled rather than a condition to be rated. Yet it retained three limitations that the present model is designed to address: a heavy dependence on expert elicitation for its key parameters, the absence of any treatment of sector-specific regulatory capital, and an underdeveloped account of how losses propagate through supply chains.

The third generation, from roughly 2015 to the present, has sought to integrate cybersecurity into the mainstream of financial risk management. Building on the foundational work that framed optimal security investment as a function of breach probability and magnitude (Gordon and Loeb, 2002; Gordon, Loeb, and Zhou, 2016), researchers conducted the first large-scale empirical analyses of cyber insurance policies, revealing inconsistent coverage driven precisely by the absence of standardized financial measurement (Romanosky, Ablon, Kuehn, and Jones, 2019). Empirical study of large samples of cyber loss events confirmed that these losses follow heavy-tailed, skewed distributions that distinguish cyber risk from conventional operational risk (Eling and Wirfs, 2019; Edwards, Hofmeyr, and Forrest, 2016; Maillart and Sornette, 2010). The annual industry studies of the cost of a data breach have provided the most widely cited empirical foundation for this generation of modelling. The present model belongs to this third generation and seeks to extend it, adding the financial denomination, sector calibration, supply-chain propagation, and regulatory cost modelling that remain incomplete in the existing literature.

2.2Structural Limitations of Existing Frameworks

A systematic review of forty-seven cybersecurity risk quantification frameworks published between 2010 and 2025 reveals five structural limitations that recur across the field and that the Financial-Based Cybersecurity Risk Exposure Model directly addresses. These limitations are not incidental shortcomings of tools; they are systematic features of how the field has approached the problem and naming them precisely is the first step toward correcting them.

2.2.1 The Financial Denomination Deficit

Of the forty-seven frameworks reviewed, thirty-eight, or roughly eighty-one percent, produce ordinal or semi-quantitative output rather than monetary estimates. This is not a cosmetic matter. Ordinal ratings cannot be aggregated across categories of risk; they cannot be integrated into regulatory capital calculations, and they cannot serve as inputs to actuarial pricing. The materiality standard applied by securities regulators, which turns on what a reasonable investor would consider important, cannot be meaningfully applied to a rating of medium-high, but it can be applied to an estimate of twelve million dollars. Earlier work demonstrated that translating qualitative cyber risk assessments into financial estimates required an entirely separate body of analytical work, confirming the gap between what the frameworks produce and what decisions require (Anderson and colleagues, 2013). The denomination deficit is, in a sense, the master limitation from which several of the others follow, because a framework that does not speak in money cannot connect to the financial machinery of the modern enterprise.

2.2.2 The Absence of Sector Calibration

Existing frameworks tend to treat all industries as equivalent for the purpose of estimating exposure. In practice, the sector to which an enterprise belongs is among the strongest determinants of the cost of a breach. Industry research has documented that the average cost of a health-care breach is several times that of a retail breach, a differential driven by the regulatory obligations attached to health information, the criticality of the services involved, and the sensitivity of the data at stake. The critical infrastructure designations maintained by the Cybersecurity and Infrastructure Security Agency reflect a federal recognition that sector membership determines systemic importance and shapes the consequences of an attack. Yet no widely used framework incorporates financial multipliers calibrated to these sector classifications, with the result that exposure estimates ignore one of the most powerful explanatory variables available.

2.2.3 The Absence of Supply-Chain Propagation

Modern cyber incidents generate losses not only at the directly affected enterprise but across networks of interconnected organizations. A single compromised software update propagated to roughly eighteen thousand organizations in one widely studied supply-chain attack. A ransomware incident at a fuel pipeline operator generated shortages across seventeen states, with the economic disruption far exceeding the direct cost of the breach. An attack on a widely used management platform propagated to fifteen hundred downstream businesses within seventy-two hours. Research has demonstrated that the correlation of cyber losses within digital supply chains significantly exceeds the assumptions embedded in existing insurance pricing models (Huang and Pearce, 2020). The component of the present model devoted to digital trade disruption addresses this gap directly, treating propagation as a measurable financial quantity rather than an unmodeled externality.

2.2.4 Inadequate Regulatory Cost Modelling

The financial consequences of cyber incidents in the United States now frequently include regulatory costs that exceed the direct expense of remediation. Data protection enforcement has generated very large cumulative fines internationally, while domestic enforcement by financial supervisors, consumer protection authorities, and state attorneys general produces substantial civil monetary penalties and mandates costly technology remediation programs. Enforcement of securities disclosure obligations relating to cybersecurity has accelerated following the implementation of the 2023 rules. Research analyzing these enforcement patterns has documented average penalties and remediation costs that are simply not captured by existing risk frameworks (Dittmar and Field, 2024). A framework that omits regulatory cost omits one of the largest and most predictable categories of financial harm.

2.2.5 Dependence on Expert Elicitation

Finally, the field has leaned heavily on the elicitation of estimates from experts, a practice shown to be systematically unreliable. Foundational work in the measurement of risk demonstrated that expert estimates tend to be overconfident within narrow ranges and underconfident about the tail events that matter most, producing assessments that are internally consistent but empirically untrustworthy (Hubbard and Seiersen, 2016). A model that anchors its key distributions to empirically derived datasets, relegating expert judgment to a secondary calibration role, can reduce this source of error substantially. The present model is built on exactly that principle.

Limitation	Consequence for decision-making	How the present model responds
Financial denomination deficit	Outputs cannot enter financial statements, capital models, or insurance pricing	All outputs are denominated in United States dollars at stated confidence levels
Absence of sector calibration	Exposure estimates ignore the strongest single predictor of breach cost	Sector Exposure Multipliers calculated to critical infrastructure classes
Absence of supply-chain propagation	Single-entity models understate exposure for connected enterprises	Digital Trade Disruption Coefficient models propagation explicitly
Inadequate regulatory cost modelling	A large and predictable category of harm is omitted	Regulatory Cost Component grounded in enforcement data
Dependence on expert elicitation	Estimates are biased and unreliable at the tail	Distributions anchored to empirical loss datasets

Table 1. The five structural limitations of existing cyber risk quantification frameworks and the corresponding design responses of the Financial-Based Cybersecurity Risk Exposure Model.

2.3The United States Financial Sector Regulatory Context

The United States financial sector provides the primary domain against which the model is calibrated, because it is in this sector that the demand for financially denominated cyber risk estimates is most acute and the regulatory expectations most explicit. The standardized approach to operational risk capital under the international bank capital accords requires banks to hold capital as a function of both an income-derived component and a historical-loss-derived component, and cyber losses feed directly into the loss component of that calculation. The stress-testing framework administered by the Federal Reserve has incorporated cyber scenarios for several years, requiring quantitative estimates of loss under specified hypothetical events. Supervisory guidance issued to national banks expects them to maintain cyber risk quantification capabilities sufficient for capital planning, stress testing, and recovery planning, an expectation that amounts to a regulatory mandate for financial-grade cyber risk models.

Beyond the domestic context, the operational resilience regime applicable to financial institutions in the European Union, in force from the start of 2025, requires financial entities and their critical technology providers to conduct financial impact assessments for scenarios of technology disruption. For United States financial institutions with European operations, this regime creates an urgent and concrete need for exactly the kind of financially denominated output that the present model produces. The convergence of these regimes means that a multinational financial institution may face several distinct regulatory demands for quantitative cyber loss estimates, each expecting output in financial terms, and a single coherent model capable of serving all of them at once offers substantial practical value.

2.4Digital Trade Infrastructure and Systemic Cyber Risk

Digital trade, by which is meant the conduct of international commerce through digital infrastructure for payment processing, invoicing, logistics coordination, and supply-chain visibility, represents a rapidly growing and under-recognized dimension of cyber risk. International economic bodies have estimated that digital trade generates economic activity measured in the tens of trillions of dollars annually, with a substantial and growing share of global goods trade now digitally intermediated. From the perspective of

cyber risk, digital trade infrastructure exhibits three properties that create distinctive exposure. It displays high network centrality in that a small number of systemically important platforms intermediate a disproportionate share of trade flows. It displays tight operational coupling, in that just-in-time logistics make even brief disruptions financially costly. And it displays jurisdictional fragmentation, in that cross-border trade involves multiple regulatory regimes with conflicting notification timelines and requirements.

Analysis by international financial institutions has documented that cyber incidents at major financial institutions generate spillover effects on trading partners that average a substantial fraction of the first-order costs, which is direct evidence that single-entity risk models systematically understate the total financial exposure of trade-infrastructure enterprises. Economic modelling has estimated that cyber-related disruptions to digital trade infrastructure could reduce national output by a measurable fraction under a severe scenario, which establishes the macroeconomic scale of the risk that goes unmodeled when propagation is ignored. These findings motivate the treatment of digital trade disruption as a first-class component of the model rather than as a secondary adjustment.

2.5Theoretical Foundations

The model draws on three established theoretical traditions and situating it within them clarifies both its logic and its limits. From financial risk theory it borrows the concepts of value-at-risk and expected shortfall, developed originally for market and credit risk, and adapts them to the cyber domain, accounting for the heavy-tailed and skewed character of cyber loss distributions that empirical research has repeatedly confirmed. From expected utility theory, it incorporates the risk aversion implicit in regulatory capital requirements, recognizing that the expected shortfall at extreme quantiles captures the systemic externalities that justify holding capital above an actuarially fair level. From collective risk theory in actuarial science, it employs the compound model in which the frequency of loss events follows a Poisson process and the severity of each loss is drawn from a heavy-tailed distribution, a structure that has been empirically validated across multiple large-scale cyber loss datasets (Klugman, Panjer, and Willmott, 2019). The synthesis of these three traditions is what allows the model to speak simultaneously to the actuary pricing a policy, the supervisor setting a capital requirement, and the board allocating a security budget.

It is worth being explicit about why the compound Poisson structure with heavy-tailed severity is the right foundation rather than a more familiar bell-shaped assumption. Cyber losses are not symmetric. Most years bring a cluster of small incidents, and occasionally a single catastrophic event dwarfs everything else. A model that assumes losses are normally distributed will badly underestimate the probability and magnitude of these rare but ruinous events, and it is precisely those events that determine capital adequacy and the sustainability of an insurance program. The log-normal and related heavy-tailed families capture this asymmetry, assigning meaningful probability to the extreme right tail. The choice of distribution is therefore not a technical nicety but a decision with direct financial consequences and grounding it in the empirical evidence is one of the ways the model resists the optimism that has undone so many earlier approaches.

2.6The Regulatory Landscape in Detail

Because the model is calibrated primarily to the United States financial sector and is intended to serve regulatory purposes, it is worth examining the regulatory landscape in greater detail than a brief mention allows. The regulatory drivers of demand for financially denominated cyber risk estimates fall into five

clusters, each of which independently strengthens the case for the kind of quantification the model provides, and which together make that quantification close to indispensable for a large enterprise.

2.6.1 Securities Disclosure

The cybersecurity disclosure rules adopted by the Securities and Exchange Commission and effective from December 2023 impose two distinct obligations. The first is incident disclosure: a registrant that experiences a cybersecurity incident must determine without unreasonable delay whether the incident is material and, if it is, must disclose the nature, scope, and timing of the incident and its material impact within four business days of the materiality determination. The second is periodic disclosure: a registrant must describe, in its annual report, its processes for assessing, identifying, and managing material risks from cybersecurity threats, together with the board's oversight of those risks and management's role in assessing and managing them. Both obligations turn on materiality, and materiality is irreducibly financial. A registrant cannot responsibly judge whether a reasonable investor would consider an incident important without an estimate, in money, of the incident's likely consequences. The model supplies precisely that estimate, and its distribution-based output allows a registrant to distinguish the typical case from the severe one when forming a materiality judgment.

2.6.2 Bank Capital

Under the international bank capital accords, banks must hold capital against operational risk, a category that explicitly includes losses from cyber incidents. The standardized approach computes the operational risk capital requirement as a function of a business indicator derived from income and a loss component derived from the bank's historical operational losses. Large cyber losses feed directly into the loss component, raising the capital requirement. A bank that can estimate its cyber loss distribution in advance can anticipate the capital consequences of a severe cyber event and plan accordingly, rather than discovering them after the fact. The model's regulatory capital adjustment sub-component captures this mechanism explicitly, computing the present value of the incremental capital cost triggered by a significant loss.

2.6.3 Operational Resilience

The operational resilience regime applicable to financial entities in the European Union, in force from the start of 2025, requires those entities and their critical technology providers to identify, assess, and manage the risks arising from disruptions to information and communication technology, including the conduct of financial impact assessments for disruption scenarios. For a United States financial institution with European operations, this regime creates a direct demand for financially denominated estimates of the consequences of technology disruption, exactly the output the model produces. The convergence of this regime with domestic stress-testing and capital requirements means that a single multinational institution may need to produce cyber loss estimates for several supervisors at once, and a model capable of serving all of them from one parameterization offers substantial efficiency.

2.6.4 Incident Notification

Banking regulators in the United States have adopted requirements that banking organizations notify their primary federal regulator of significant computer-security incidents within a short window, typically thirty-six hours of determining that a notification incident has occurred. While these requirements are procedural rather than quantitative, they sharpen the need for an enterprise to understand, quickly and in financial terms, the likely consequences of an incident, because the same incident that triggers a notification obligation will often trigger a materiality assessment and a capital consequence. An enterprise that has modelled its exposure in advance is better placed to respond coherently across all of these obligations under the time pressure that a real incident imposes.

2.6.5 State Privacy and Sectoral Regimes

Beyond the federal financial regulators, a growing patchwork of state privacy statutes and sectoral regimes governing health information, consumer financial data, and other categories creates additional exposure to civil monetary penalties and mandated remediation. The decision-tree approach embedded in the model's regulatory cost component is designed to accommodate this patchwork, mapping the characteristics of a breach to the regimes that apply and estimating the resulting penalty exposure. As the patchwork grows, the value of a model that can integrate exposure across many regimes into a single financial estimate grows with it.

2.7The Economics of Cyber Risk and Security Investment

The model rests on an economic logic that deserves explicit statement, because that logic is what connects security control to a financial outcome. The foundational economic treatment of security investment frames the problem as one of optimal expenditure: an enterprise should invest in security up to the point at which the marginal reduction in expected loss equals the marginal cost of the investment. This frame is intuitive, but it has long been frustrated by the absence of a credible estimate of expected loss in financial terms. Without such an estimate, the marginal reduction in loss cannot be computed, and the optimization collapses into guesswork. The model resolves this frustration by producing the financial loss estimate the optimization requires, and by allowing the marginal effect of a control to be computed directly through the device of re-running the model with and without control.

A second economic consideration concerns the difference between expected loss and the loss that matters for capital and solvency. An enterprise that provisioned only against its expected annual loss would be undercapitalized against the severe events that determine survival, because the distribution of cyber losses is heavily tailed and the severe event is far larger than the average. This is why the model reports not only the median but the ninetieth and ninety-ninth percentiles and the expected shortfall beyond the tail. The economics of solvency are governed by the tail, not by the mean, and a model that reported only the mean would mislead precisely where the stakes are highest. The risk aversion implicit in regulatory capital requirements is an expression of this same logic: supervisors require capital above the actuarially fair level precisely because the social cost of a financial institution's failure exceeds the private cost, and the tail of the loss distribution is where that divergence lives.

2.8The Empirical Study of Cyber Loss Distributions

A substantial body of empirical research has examined the statistical character of cyber losses, and the model's distributional choices are grounded in that research. Studies of large samples of cyber loss events have repeatedly found that the distribution of loss magnitudes is heavy-tailed and positively skewed, meaning that most losses are modest, but a small number are very large, and that the very large losses contribute a disproportionate share of the total. This finding distinguishes cyber risk from many conventional risks whose losses cluster more tightly around a central value. The log-normal distribution, and related heavy-tailed families, capture this character well, which is why the model adopts a log-normal form for the loss magnitude term. Research has also examined the frequency of cyber events and found it consistent with a Poisson process at the level of the individual enterprise, which justifies the compound Poisson structure the model employs. The combination of Poisson frequency and a heavy-tailed severity is the workhorse model of actuarial science for exactly this kind of risk, and its application to cyber risk is both natural and empirically supported.

It is worth noting the tension in this literature that the model must navigate. The heavy tail of the cyber loss distribution means that the historical record, however large, may not contain an event as severe as the worst event that could occur, simply because the worst events are rare and the record is finite. This is the problem of the unseen tail, and it means that any model calibrated to historical data risks understating the true severity of the most extreme outcomes. The model addresses this in two ways: by adopting a distributional form whose tail extends beyond the observed maximum, assigning meaningful probability to events more severe than any yet recorded, and by treating nation-state events, which include some of the most severe outcomes ever observed, through a separate term with an even heavier tail. The honest acknowledgement of the unseen tail is part of what distinguishes a sober risk model from an exercise in false precision.

2.9The Gap in Digital Trade Literature

A distinct gap in existing literature concerns the intersection of cyber risk and digital trade, and closing it is among the paper's contributions. The literature on digital trade has documented its scale and growth, establishing that digitally intermediated commerce now accounts for a large and rising share of global economic activity. The literature on cyber risk has documented the heavy-tailed character of cyber losses and the propagation of losses through supply chains. But the two literatures have largely developed in parallel, and the specific question of how cyber risk concentrates within and propagates through digital trade infrastructure has received comparatively little quantitative treatment. This is a consequential omission, because digital trade infrastructure exhibits exactly the properties, high centrality, tight coupling, and jurisdictional fragmentation, that make cyber risk more dangerous and most likely to propagate. The trade disruption component of the model is, in part, an attempt to bring these two literatures together, treating the propagation of cyber losses through digital trade networks as a measurable financial quantity and thereby filling a gap that neither literature alone has filled. The validation finding that propagation dominates the exposure of the most trade-intensive enterprises is direct evidence that this gap matters, and that a framework which ignores it will misjudge the exposure of precisely the enterprises on which the digital economy most depends.

2.10 Milestones in Intellectual History

The three-generation account of the field can be enriched by naming the milestones that marked the passage from one generation to the next, because those milestones reveal the logic of the field's development. The first generation was defined by the codification of information security management into systematic

standards and by the first federal methodologies for assessing risk, achievements that established cybersecurity as a discipline with structured processes but that expressed risk in qualitative terms. The transition to the second generation was marked by the introduction of standardized technical severity scoring, which gave the field a common vocabulary for describing the danger of a vulnerability, and above all by the development of a model that decomposed cyber risk into frequency and magnitude and advocated simulation over distributions, the first widely adopted attempt to treat cyber risk as a quantity rather than a condition.

The transition to the third generation was marked by the integration of cybersecurity into financial economics, beginning with the formalization of optimal security investment as a function of breach probability and magnitude, continuing through the first large-scale empirical analyses of cyber insurance and cyber losses, which established the heavy-tailed character of the loss distribution, and arriving at the present moment, in which regulatory mandates for financially denominated disclosure and capital have made the financial measurement of cyber risk not merely an academic aspiration but a practical necessity. The model developed in this paper is best understood as an attempt to consummate this third generation, to carry the integration of cybersecurity into financial risk management to the point at which a single auditable equation produces a monetary estimate suitable for the balance sheet, the capital calculation, the insurance contract, and the regulatory filing. Each generation built on the achievements and confronted the limitations of its predecessor, and the present model stands on the accumulated work of all three.

2.12 A Brief Comparison with International Approaches

Although the model is calibrated to the United States, it is worth situating it briefly against the approaches taken in other jurisdictions, because the comparison clarifies both what is general in the model and what is specific to its calibration. The European approach to cyber risk in the financial sector, embodied in the operational resilience regime, places heavy emphasis on the assessment of financial impact under disruption scenarios, an emphasis that aligns closely with the model's output and that makes the model directly applicable to European compliance once its regulatory cost component is recalibrated to the European enforcement landscape. The international bank capital accords, which apply across many jurisdictions, provide a common framework for operational risk capital into which the model's estimates feed, regardless of the specific national supervisor. International standard-setting bodies have promoted common practices for cyber incident response and recovery and for the resilience of financial market infrastructure, and the model is consistent with the spirit of those practices, providing the financial quantification they increasingly assume. The structure of the model, its four components, and its simulation architecture are general and would carry across jurisdictions; what is specific to the United States is the calibration of the regulatory cost component and the sector multipliers, both of which a practitioner in another jurisdiction would recalibrate to local conditions. The model is thus best understood as a general methodology with a United States calibration, and its extension to other jurisdictions is a matter of recalibration rather than redesign.

2.13 The Relationship Between Cyber Risk and Operational Risk

A conceptual question that underlies much of the model concerns the relationship between cyber risk and the broader category of operational risk, and clarifying it sharpens the model's contribution. In the taxonomy of bank supervision, cyber risk is a subset of operational risk, the risk of loss from inadequate or failed internal processes, people, and systems, or from external events. This classification has the advantage of bringing cyber losses within the established apparatus of operational risk capital, but it has the disadvantage

of obscuring the features that make cyber risk distinctive. Cyber losses are more heavy-tailed than typical operational losses, more prone to propagate across networks, more entangled with regulatory consequences, and more sensitive to the deliberate adaptation of an intelligent adversary, who, unlike a flood or a process failure, changes tactics in response to defenses. Treating cyber risk merely as ordinary operational risk and modelling it with tools designed for operational losses generally understates these distinctive features and produces the systematic errors the benchmarks in this paper display. The model's contribution is to honor the classification by producing outputs that feed the operational risk capital calculation while respecting the distinctiveness, by modelling the heavy tail, the propagation, the regulatory entanglement, and the adversarial character that generic operational risk tools miss. Cyber risk is operational risk, but it is operational risk of a particular and dangerous kind, and a model that recognizes both half of that statement serves the practitioner better than one that recognizes only the first.

3. The Model Framework: Mathematical Specification

This section specifies the model in full. The exposition proceeds from the overall architecture to each of the four components in turn, and then to the form of the output and the reporting standards that accompany it. Throughout, the aim is not only to state the equations but to explain the reasoning behind each modelling choice, so that a practitioner can both implement the model and defend its assumptions to a board, an auditor, or a supervisor.

3.1Model Architecture and Core Equation

The model produces a Cyber Risk Exposure Score, expressed in United States dollars, by integrating four analytically distinct components, each of which captures a different financial dimension of an enterprise's cyber exposure. The first is the Annualized Loss Expectancy, which estimates the direct cost of breaches. The second is the Sector Exposure Multiplier, which adjusts that direct cost for the systemic and regulatory consequences that vary by industry. The third is the Digital Trade Disruption Coefficient, which captures the secondary losses that propagate through trade networks. The fourth is the Regulatory Cost Component, which estimates the financial obligations arising from the regulatory response to an incident. The core equation that combines them is as follows.

$$\mathbf{CRES = (ALE \times SEM) + (DTDC \times TDE) + RC}$$

Here the Cyber Risk Exposure Score, abbreviated as the score in the discussion that follows, is the sum of three terms: the product of the Annualized Loss Expectancy and the Sector Exposure Multiplier, which represents direct loss adjusted for sector consequence; the product of the Digital

Trade Disruption Coefficient and the Total Digital Trade Exposure, which represents propagated trade losses; and the Regulatory Cost Component, which represents regulatory exposure. The additive structure is deliberate. It allows each term to be inspected, audited, and challenged independently, and it allows a practitioner to see briefly which dimension of exposure dominates a given enterprise's profile. A bank with modest trade dependencies but heavy regulatory exposure will show a different composition from a digital trade platform whose propagated losses dwarf its direct ones, and the model makes that difference visible rather than burying it in a single opaque number.

Figure 2. Model architecture: data flow and component integration. Input data sources feed four analytical components whose outputs are combined through the core equation to produce a probability-distribution-based monetary risk score. [Figure to be inserted from the original manuscript.]

3.2Annualized Loss Expectancy

The Annualized Loss Expectancy is computed through a two-stage Monte Carlo simulation across ten thousand iterations. In the first stage, the annual frequency of breaches is modelled as a Poisson process with a rate calibrated from three empirical sources: the industry-sector baseline frequency reported in the annual data breach investigations literature, an adjustment for enterprise size based on documented relationships between size and attack surface, and an adjustment for control maturity derived from an assessment of implemented controls against an established controls catalog. In the second stage, the magnitude of each breach loss is drawn from a sector-calibrated heavy-tailed distribution whose parameters are fitted against an empirical cyber loss database, segmented by sector and by revenue cohort. The expanded form of the component is given below.

$$\mathbf{ALE = TEF \times V \times (1 - CS) \times LM\ (\mu, \sigma)}$$

In this expression, the Threat Event Frequency is the annualized Poisson rate, calibrated from breach frequency data and threat intelligence. The vulnerability coefficient, bounded between zero and one, represents the probability that a threat event produces a loss event given the controls currently in place, calibrated against documented technique success rates by sector. The control strength factor, also bounded between zero and one, represents the proportional reduction in loss probability achieved by implemented controls, scored against an established controls catalog. The loss magnitude term is a heavy-tailed distribution whose location and scale parameters are fitted to sector-specific empirical loss data.

The selection of a log-normal form for the loss magnitude distribution is supported by empirical validation across several independent datasets, by its established fit to operational risk losses more generally, and by its tractable tail behavior, which permits reliable estimation of extreme quantiles. From the distribution of simulated annual losses, the model extracts decision-relevant quantiles. The value at the ninetieth percentile serves as the input to the core equation for purposes of risk appetite and insurance, while the value at the ninety-ninth percentile is used for regulatory stress testing, where the concern is precisely with severe but plausible events rather than with the typical year.

3.3The Sector Exposure Multiplier

The Sector Exposure Multiplier adjusts the Annualized Loss Expectancy for the differential financial consequences of cyber incidents across the critical infrastructure sectors. It reflects four sector-specific factors: the severity of regulatory penalties under the applicable legal regimes; the degree of systemic economic interconnectedness, derived from national input-output accounts; the value of the sector's data assets, reflecting the price such data commands in illicit markets; and the cost of recovery time, reflecting the financial value of service continuity in the sector's operational context. The multiplier is computed from these four sub-components through a weighted linear combination, with weights derived from a structured pairwise comparison procedure conducted during expert calibration.

$$SEM = 0.35 \times Reg + 0.30 \times Syst + 0.20 \times Data + 0.15 \times RTO$$

The allocation of weights reflects the relative financial materiality of each sub-component as assessed by the expert panel and validated against the loss dataset. Regulatory penalty severity receives the highest weight, at thirty-five percent, because enforcement actions generate the largest and most predictable financial obligations following a breach and are directly observable from public records. Systemic interconnectedness receives the second-highest weight, at thirty percent, reflecting the empirical finding that systemic contagion effects account for a substantial share of total economic loss in high-centrality sectors. Data asset value receives twenty percent, reflecting the market price of sector-specific data in illicit marketplaces. Recovery time cost receives the lowest weight, at fifteen percent, because, while highly sector-specific, it primarily affects the business interruption component already partially captured in the Annualized Loss Expectancy rather than the regulatory and systemic dimensions that differentiate sectors most sharply. Sensitivity testing confirmed that variations of five percentage points around these weights produce less than eight hundredths of a unit of deviation in the multiplier across all sectors, indicating that the composite is robust to modest misspecification of the weights.

Sector	SEM	Reg	Syst	Data	RTO $/hr	Primary financial driver
Financial Services	3.2	3.5	3.8	3.4	$1.2M	Systemic contagion; bank capital surcharge
Healthcare and Public Health	2.9	2.8	2.4	3.9	$0.9M	Health data value; health-privacy enforcement
Energy	3.1	2.4	3.9	2.6	$1.5M	Grid interdependency; control-system cascades
Information Technology	2.7	2.0	3.5	3.1	$0.7M	Software supply-chain multiplier
Transportation Systems	2.5	1.8	3.2	2.3	$0.8M	Just-in-time logistics; trade continuity
Communications	2.8	2.3	3.6	2.7	$0.6M	Network backbone; cascading failure risk
Défense Industrial Base	2.6	1.9	2.8	3.3	$0.9M	Intellectual property loss; contract penalties
Water and Wastewater	2.3	2.5	3.1	1.8	$0.4M	Public health consequences; environmental enforcement
Food and Agriculture	1.9	1.5	2.4	1.7	$0.3M	Supply-chain perishability; food-safety oversight
Other sectors (average)	1.8	1.4	2.0	1.9	$0.2M	Sector-specific regulatory adjustment

Table 2. Sector Exposure Multipliers by critical infrastructure sector, with component sub-scores and primary financial drivers. Sub-scores are on a scale from one to four; the multiplier is a composite weighted value. Financial Services and Energy carry the highest multipliers, reflecting systemic interconnectedness, regulatory penalty severity, and high recovery-time costs.

The calibration of these values followed a multi-stage process: a structured expert assessment across twelve sector specialists drawn from public-sector cybersecurity bodies, financial-sector information-sharing organizations, and corporate security functions; validation against ten years of breach cost data segmented by sector; and iterative adjustment to minimize estimation error against the validation dataset. The result is a set of multipliers that reproduces the documented differentials in breach cost across sectors, with financial services and energy carrying the highest values because of their systemic interconnectedness and the severity of the regulatory and operational consequences they face.

3.4The Digital Trade Disruption Coefficient

The Digital Trade Disruption Coefficient captures the secondary financial exposure that arises when a cyber incident propagates through digital trade networks, disrupting supply chains, interrupting payment flows, triggering contractual penalties, and eroding correspondent relationships. It is the most novel component of the model and addresses the largest documented gap in existing frameworks. Its form is given below.

$$\mathbf{DTDC = DTDR \times NCS \times (1 + HPV / 30)}$$

The digital trade dependency ratio, bounded between zero and one, is the proportion of total trade value intermediated through digital infrastructure, derived from enterprise financial statements and an analysis of digital channels. The network centrality score, also bounded between zero and one, is the normalized centrality of the enterprise within its mapped digital trade network, reflecting its systemic importance to the continuity of that network. The historical propagation velocity, measured in days, is the empirically observed time to impact on trade partners following comparable breach events; faster propagation, meaning a lower number of days, yields a higher coefficient, reflecting greater acute financial exposure. The total digital trade exposure against which the coefficient is applied is derived from enterprise financial data, namely the revenue, payables, receivables, and inventory value attributable to digitally intermediated trade relationships, adjusted to capture only the digitally dependent portions. The product of the coefficient and the exposure yield a monetary estimate of secondary trade disruption losses, additive to the direct losses captured in the first term of the core equation.

3.5The Regulatory Cost Component

The Regulatory Cost Component models the financial obligations that arise specifically from the regulatory response to a cyber incident in the United States. It comprises three sub-components, summed as shown.

$$RC = CMP + MRE + RCA$$

3.5.1 Civil Monetary Penalties

The first sub-component estimates expected regulatory fine exposure across the applicable domestic and international regimes. A decision-tree model maps the characteristics of a breach, including the categories of data involved, the volume of records compromised, the jurisdiction of the affected data subjects, the prior regulatory history of the enterprise, the timeliness of notification, and the quality of controls in place before the incident, to expected penalty ranges under the relevant statutes governing data protection, consumer privacy, health information, financial supervision, and securities disclosure. The estimate is the probability-weighted average across enforcement outcome scenarios, calibrated against six hundred and eighty publicly disclosed United States regulatory enforcement actions from 2015 to 2025, covering banking supervisors,

the securities regulator, the health-privacy enforcement authority, the consumer protection authority, and state attorneys general.

3.5.2 Mandated Remediation Expenditure

The second sub-component captures the costs of the technology investment, process redesign, independent audit, enhanced monitoring, and compliance with consent decrees that regulators mandate following significant incidents. It is calibrated against the same enforcement dataset, with consent orders involving systemic control failures generating estimates several times higher than incident-specific remediation, reflecting the comprehensive technology transformation programs typical of the most serious supervisory responses.

3.5.3 Regulatory Capital Adjustment

The third sub-component applies to financial institutions subject to the International Bank Capital Accords. For these institutions, a large cyber loss increases the loss component of the operational risk capital calculation, raising the minimum capital requirement. The adjustment is computed as the present value of the incremental cost of capital over the expected supervisory response period, using a cost of equity estimated through a standard asset pricing approach. For enterprises not subject to these requirements, this sub-component is zero.

3.6Output and Reporting Standards

The model produces a probability distribution over possible annual cyber loss values, from which practitioners extract the metrics relevant to their decisions. The median is used for budget planning. The ninetieth percentile is used for setting risk appetite and determining insurance limits. The ninety-ninth percentile is used for regulatory stress testing. The conditional expectation beyond the tail, sometimes called expected shortfall, is used for reinsurance design and systemic risk analysis. Results are accompanied by confidence intervals derived from the simulation, by a coefficient of variation that conveys the relative uncertainty of the estimate, and by a ranking of the five model inputs with the greatest influence on the result, so that the practitioner knows where to concentrate effort to improve precision. This discipline of reporting not only a point estimate but a distribution, an uncertainty measure, and a sensitivity ranking is part of what distinguishes a financially credible model from a single number presented without context.

3.7Calibration Methodology in Depth

The credibility of the model rests on the quality of its calibration, and so the calibration procedure deserves a fuller account than the component specifications above provided. Calibration proceeds in four stages, each of which produces parameters that feed the next.

The first stage calibrates the frequency parameters. The annualized threat event frequency for a given enterprise begins with a sector baseline derived from breach-investigation data, which establishes the typical rate of loss events for enterprises of the relevant sector. This baseline is then adjusted for the enterprise's size, on the principle that a larger attack surface attracts more attempts, and for its control maturity, on the principle that stronger controls convert fewer attempts into losses. The size adjustment is

derived from documented relationships between enterprise scale and incident frequency; the control adjustment is derived from an assessment of the enterprise's controls against an established catalog, expressed as the control strength factor in the loss equation. The output of the first stage is an enterprise-specific frequency parameter.

The second stage calibrates the severity parameters. The location and scale parameters of the log-normal loss magnitude distribution are fitted to sector-specific empirical loss data, segmented by revenue cohort, so that the distribution reflects the scale of the enterprise as well as its sector. Where enterprise-specific loss data is available, it is blended with the sector data to sharpen the fit; where it is not, the distribution sector is used directly. The output of the second stage is an enterprise-specific severity distribution.

The third stage calibrates the sector multiplier and its sub-components through the structured expert procedure described earlier, validated against a decade of sector-segmented breach cost data. The fourth stage calibrates the regulatory cost component through the decision-tree model fitted to the enforcement dataset. The four stages together produce a complete parameterization, and the entire parameterization is then validated against the holdout of realized losses to confirm that the assembled model reproduces the historical record within the stated error.

3.8The Monte Carlo Procedure

The two-stage Monte Carlo simulation at the heart of the model warrants explicit description, because it is the engine that converts parameters into a distribution of outcomes. The procedure is as follows, and it is run for ten thousand iterations to produce the loss distribution and for fifty thousand iterations when the sensitivity analysis requires a finer resolution of the tail.

1. Draw the annual count of loss events from a Poisson distribution whose rate is the calibrated threat event frequency, adjusted by the vulnerability coefficient and the control strength factor.

2. For each loss event in the drawn count, draw a loss magnitude from the calibrated log-normal severity distribution.

3. Sum the loss magnitudes across all events in the iteration to obtain the total direct loss for that simulated year.

4. Apply the sector multiplier to the direct loss, add the trade disruption term computed from the enterprise's dependency, centrality, and exposure, and add the regulatory cost term computed from the decision-tree model.

5. Record the resulting score for the iteration.

After all iterations are complete, the recorded scores form an empirical distribution from which the median, the ninetieth and ninety-ninth percentiles, and the expected shortfall beyond the tail are extracted, and from which confidence intervals are derived by resampling. The number of iterations is chosen to make the estimates of the extreme quantiles stable; ten thousand iterations is sufficient for the ninetieth percentile and adequate for the ninety-ninth, while the larger count used in the sensitivity analysis improves the resolution of the variance decomposition at the tail. The procedure is deliberately simple to state, because a model intended for regulatory and audit use must be reproducible by a third party, and a procedure that can be described in five steps can be audited in five steps.

3.9Notation and Assumptions

For reference, the principal symbols of the model and their meanings are collected below, followed by a statement of the model's central assumptions.

Symbol	Meaning
CRES	Cyber Risk Exposure Score is the monetary output of the model at a stated confidence level
ALE	Annualized Loss Expectancy is the direct loss term derived from a Monte Carlo simulation
SEM	The Sector Exposure Multiplier is the composite adjustment for sector consequence
DTDC	Digital Trade Disruption Coefficient, capturing propagation through trade networks
TDE	Total Digital Trade Exposure, the digitally intermediated portion of enterprise trade value
RC	Regulatory Cost Component, the sum of penalties, remediation, and capital adjustment
TEF	Threat Event Frequency, the annualized Poisson rate of loss events
CS	Control Strength factor, the proportional reduction in loss probability from controls
NCS	Network Centrality Score, the enterprise's systemic importance in its trade network
HPV	Historical Propagation Velocity, the observed time to impact on trade partners

Table 9. Principal notation of the model.

The model rests on four central assumptions, each of which is defensible, but each of which a practitioner should hold consciously. First, it assumes that the frequency of loss events follows a Poisson process at the level of the individual enterprise, an assumption supported by empirical literature. Second, it assumes that loss magnitudes follow a heavy-tailed distribution of the log-normal family, again supported by literature. Third, it assumes that the distributional parameters are stationary over the assessment horizon, an assumption that the rapid evolution of the threat landscape strains, and that the requirement of annual recalibration is designed to mitigate. Fourth, it assumes that the four components of the score are additive and can be estimated independently, an assumption that simplifies auditing and that the additive structure of the core equation reflects. Where any of these assumptions is materially violated for a particular enterprise, the practitioner should adjust the model accordingly rather than apply it mechanically.

3.10 The Digital Trade Disruption Coefficient in Depth

Because the trade disruption component is the model's most novel element and the one validation shows to be most consequential for trade-intensive enterprises, it warrants a fuller treatment of its three factors. The dependency ratio measures how much of the enterprise's trade value flows through digital infrastructure, and it is computed from the enterprise's financial statements by attributing revenue, payables, receivables, and inventory to digitally intermediate relationships. An enterprise whose trade is almost entirely digital, such as a pure digital trade platform, has a ratio approaching one; an enterprise whose trade is largely conducted through traditional channels has a much lower ratio. The centrality score measures the enterprise's systemic importance within its trade network, computed as a normalized centrality of the enterprise within a graph of its trade relationships. An enterprise that many others depend upon, occupying

a hub position in the network, has a high centrality score and would propagate disruption widely if compromised; a peripheral enterprise has a low score. The propagation velocity measures how quickly disruption reaches trade partners after a comparable incident, expressed in days, with faster propagation producing a larger coefficient because it leaves less time to contain the financial harm.

The interaction of these three factors is what gives the component its explanatory power. An enterprise can have high dependency but low centrality, in which case its own operations are exposed, but its compromise would not propagate widely. It can have high centrality but moderate dependency, in which case it is a critical node whose compromise would ripple through the network even if its own trade is partly traditional. The most dangerous combination, from a systemic perspective, is high dependency and high centrality together, the profile of the digital trade platform archetype, because such an enterprise is both heavily exposed to itself and capable of propagating that exposure across the network it anchors. The component captures all these cases through the multiplication of its factors, and the validation confirms that this structure reproduces the observed differences in exposure across enterprises of different network positions.

3.11 The Regulatory Cost Decision Tree

The regulatory cost component's civil monetary penalty sub-component is computed through a decision tree, and the logic of that tree deserves description because it is what converts the characteristics of a breach into an expected regulatory cost. The tree begins by classifying the categories of data involved, since the regime that applies and the severity of the penalty depend heavily on whether the data is health information, consumer financial information, personal information governed by state statutes, or material non-public information governed by securities law. It then branches on the volume of records compromised, since penalties often scale with the number of affected individuals. It branches further on the jurisdiction of the affected data subjects, since cross-border breaches invoke multiple regimes with different penalty structures. It branches on the enterprise's prior regulatory history, since repeat violations attract heavier penalties. It branches on the timeliness of notification, since delayed notification is itself frequently penalized. And it branches on the quality of controls in place before the incident, since regulators distinguish between enterprises that suffered a breach despite reasonable controls and those whose controls were deficient. Each path through the tree terminates in an expected penalty range, calibrated from the enforcement dataset, and the sub-component is the probability-weighted average across the plausible paths. This structure makes the regulatory cost estimate auditable, because an examiner can trace the path the model took and verify each branch against the facts of the enterprise.

3.12 The Sector Multiplier Weighting Procedure

The weights that combine the four sub-components of the multiplier sector were derived through a structured pairwise comparison procedure during the expert calibration, and the rationale for the resulting allocation merits a fuller explanation than the headline weights convey. In a pairwise comparison procedure, experts judge the relative importance of each pair of factors, and the judgments are synthesized into a consistent set of weights. The procedure produced the allocation of thirty-five percent to regulatory penalty severity, thirty percent to systemic interconnectedness, twenty percent to data asset value, and fifteen percent to recovery time cost. Regulatory severity received the highest weight because enforcement actions are the largest and most predictable financial consequence of a breach and are directly observable from public records, which makes them both material and reliable. Systemic interconnectedness received the second weight because the empirical evidence on contagion shows it to account for a large share of total

economic loss in high-centrality sectors. Data asset value received a moderate weight reflecting the illicit market price of sector-specific data. Recovery time cost received the lowest weight, not because it is unimportant, but because much of its effect is already captured in the direct-loss term, so giving it a high weight in the multiplier would double-count it. The robustness of this allocation was confirmed by sensitivity testing, which showed that modest variations in the weights produce only small changes in the multiplier, so the model's outputs do not hinge on the precise values of weights that are, after all, the product of expert judgment.

3.13 Correlation and Dependence Between Components

The additive structure of the core equation, in which the direct-loss, trade-disruption, and regulatory-cost terms are summed, invites a methodological question that a careful practitioner should consider: does the additivity assume that the three terms are independent, and if they are in fact correlated, does the model misstate the total? The question is a fair one because the same severe incident that generates a large direct loss is also likely to generate substantial propagation and a heavy regulatory response, so the terms are positively correlated across incidents rather than independent. The model addresses this in two ways. First, the terms are not estimated from independent data but from the same underlying incidents, so the empirical calibration already reflects the joint behavior of the terms rather than treating them as unrelated. Second, because the terms are positively correlated, summing them at a common confidence level, such as the ninetieth percentile, is conservative for the typical case and appropriate for risk appetite purposes, because it reflects the reality that a bad year is bad across all three dimensions at once.

Where a practitioner requires a more precise treatment of joint distribution, for example, in designing a reinsurance program sensitive to the correlation of extreme losses, the model can be extended to represent the dependence between the terms explicitly through a joint simulation rather than a sum of marginal quantiles. This refinement is most valuable at the extreme tail, where the correlation of the terms most affects the result, and it is identified as a candidate for future development. For the great majority of applications, the additive structure, calibrated against jointly observed incidents, provides an estimate that is both auditable and appropriately conservative, and the explicit modelling of dependence is a refinement to be reached for when the decision at hand turns on the precise behavior of the joint tail.

3.14 Confidence Intervals and the Bootstrap

The model reports not only point estimates of its quantiles but confidence intervals around them, and the method by which those intervals are derived deserves description, because the intervals are central to the honest communication of uncertainty. The intervals are derived by the bootstrap, a resampling technique in which the simulated loss distribution is repeatedly resampled with replacement, the quantile of interest is recomputed for each resample, and the spread of the recomputed quantiles defines the interval. The bootstrap is well-suited to this task because it does not assume the shape of the sampling distribution of the quantile, which is valuable when the underlying loss distribution is heavy-tailed, and the quantile of interest lies in the tail, where parametric approximations are least reliable. The width of the resulting interval conveys how much confidence the practitioner should place in the point estimate: a narrow interval indicates a well-determined estimate, while a wide interval, common at the extreme tale where data is sparse, signals

that the estimate should be treated with caution. Reporting the interval alongside the point estimate is what allows a board or an auditor to weigh the estimate appropriately, and it is part of the discipline that distinguishes a credible quantitative model from one that projects a false precision. An estimate without an interval invites confidence that the data may not support; an estimate with an interval invites the considered judgment that the situation requires.

4. Worked Numerical Examples

A model is only as useful as it is usable, and the surest way to make a quantitative framework usable is to show the arithmetic. This section traces the calculation of the Cyber Risk Exposure Score for each of the three enterprise archetypes used in validation, moving step by step from the raw inputs to the final score. The inputs shown are representative values consistent with each archetype's profile, chosen to illustrate the mechanics of the model rather than to report new empirical measurements; the realized validation results are presented separately in Section 5. The purpose here is pedagogical: a risk analyst who follows these examples should be able to reproduce the method on an enterprise of their own.

> **How to read these examples**
>
> Each example proceeds in four steps. Step one computes the Annualized Loss Expectancy from the frequency, vulnerability, control, and loss-magnitude inputs. Step two applies the Sector Exposure Multiplier. Step three computes the digital trade disruption term. Step four adds the regulatory cost term and reports the combined score at the ninetieth percentile. All monetary values are in United States dollars.

4.1Archetype A: A Mid-Tier Commercial Bank

The first archetype is a nationally chartered commercial bank with total assets of roughly eighteen billion dollars, two hundred and eighty branches, and a business mix spanning retail and commercial banking. It is classified in the financial services sector. It maintains connectivity to the international payment messaging network and a correspondent banking network, which gives it a digital trade dependency ratio of 0.42 and a network centrality score of 0.31. It is subject to oversight by the principal national banking supervisors.

In the first step, the Annualized Loss Expectancy is computed. Suppose the simulation is parameterized with a threat event frequency reflecting the banking sector baseline, a vulnerability coefficient reduced by a mature control environment, and a loss magnitude distribution fitted to financial-sector breaches. The two-stage Monte Carlo simulation across ten thousand iterations produces a distribution of annual direct losses whose ninetieth-percentile value, used for risk appetite purposes, is approximately one and one-tenth million dollars before sector adjustment. In the second step, this figure is multiplied by the Sector Exposure Multiplier for financial services, which is 3.2, reflecting the systemic interconnectedness and regulatory severity of the sector. The adjusted direct-loss term is therefore approximately three and one-half million dollars.

In the third step, the digital trade disruption term is computed. With a dependency ratio of 0.42 and a centrality score of 0.31, and a historical propagation velocity consistent with the banking sector, the

coefficient is modest and applied to the bank's digitally intermediated trade exposure; it contributes a smaller secondary loss than it would for a trade-intensive enterprise. In the fourth step, the regulatory cost term is added, reflecting the bank's exposure to civil monetary penalties, mandated remediation, and the capital adjustment that applies to institutions subject to the bank capital accords. Combining the three terms yields a Cyber Risk Exposure Score at the ninetieth percentile of approximately four and two-tenths million dollars. The composition is instructive: the direct-loss term dominates, the regulatory term is material, and the trade term, while present, is the smallest of the three. This is the characteristic profile of a traditional deposit-taking institution whose exposure is concentrated in its own systems and its regulatory obligations rather than in propagation across a trade network.

Step and term	Inputs (representative)	Contribution to score
Direct loss (ALE × SEM)	ALE at ninetieth percentile near $1.1M; sector multiplier 3.2	Approximately $3.5M (the largest term)
Trade disruption (DTDC × TDE)	Dependency 0.42; centrality 0.31	Modest secondary loss
Regulatory cost (RC)	Penalties, remediation, capital adjustment	Material
Score at ninetieth percentile	Sum of the three terms	Approximately $4.2M

Table 3. Worked example for the commercial bank archetype, illustrating how the three terms of the core equation combine. Values are representative and chosen to illustrate the method.

4.2Archetype B: A Cross-Border Digital Trade Platform

The second archetype is a digital trade finance intermediary incorporated in the United States, processing roughly two and four-tenths billion dollars in annual trade flows for some twelve hundred small and medium importers and exporters across North America and Europe. It is classified in the information technology sector for purposes of the multiplier sector, but its defining feature is its extraordinary dependence on digital trade, with a dependency ratio of 0.96 and a high network centrality score of 0.78. It is subject to oversight as a money transmitter and by consumer-financial and financial-crime authorities.

The contrast with the bank is illuminating. The direct loss term, computed as before and adjusted by the information technology sector multiplier of 2.7, is meaningful but not dominant. What dominates is the trade disruption term. With a dependency ratio close to one and a centrality score of 0.78, the coefficient is large and applied to the platform's very high digitally intermediated exposure, it generates a secondary loss that constitutes the single largest component of the score. The regulatory term is present but smaller in relative terms than for the bank, because the platform is not subject to the bank capital accords in the same way. Combining the terms yields a Cyber Risk Exposure Score at the ninetieth percentile of approximately eleven and eight-tenths million dollars. For this enterprise, ignoring supply-chain propagation, as a single-entity model would, would understate the true exposure by a large margin. This is precisely the failure mode the model is built to correct.

4.3Archetype C: A Financial Technology Hybrid Operator

The third archetype is an embedded payments and consumer lending platform incorporated in the United States and the European Union, serving some four and two-tenths million customers. It straddles two sectors, financial services and information technology, and so requires a revenue-weighted blend of the two sector multipliers. Its dependency ratio is 0.71 and its centrality score is 0.54, placing it between the bank and the pure trade platform. It is subject to a wide range of domestic and European oversight.

As one would expect from its intermediate position, the financial technology hybrid shows a more balanced composition than either of the other two archetypes. The direct-loss term, adjusted by the blended multiplier, is the largest single component, but the trade disruption term is substantial, and the regulatory term is material, given the breadth of the regimes to which the operator is subject. Combining the terms yields a Cyber Risk Exposure Score at the ninetieth percentile of approximately seven and one-tenth million dollars. The lesson of the three examples taken together is that the same equation produces very different compositions for very different enterprises, and that the composition itself, not merely the headline number, is a source of insight for the risk practitioner.

Term	Bank (A)	Trade platform (B)	Fintech hybrid (C)
Direct loss share	59.6 percent	43.5 percent	53.0 percent
Trade disruption share	18.3 percent	41.7 percent	27.4 percent
Regulatory cost share	22.1 percent	14.8 percent	19.6 percent
Score at the 90th percentile	$4.2M	$11.8M	$7.1M

Table 4. Component composition of the Cyber Risk Exposure Score across the three archetypes. The trade disruption share rises sharply for the trade-intensive platform, while direct loss dominates for the traditional bank.

4.4A Sensitivity Walkthrough

The worked examples gain further value when one input is varied, and the effect on the score is traced, because this is exactly the exercise a practitioner performs when evaluating a control investment or testing the robustness of a result. Consider the commercial bank archetype and suppose its control strength improves, as it would if the bank deployed enhanced monitoring and identity controls. The improvement reduces the probability that a threat event becomes a loss event, which lowers the direct-loss term of the score. Because the direct-loss term is the largest component for the bank, accounting for nearly sixty percent of its score, an improvement in control strength produces a meaningful reduction in the total, and the magnitude of that reduction is exactly the financial return on the control investment that the enterprise risk application uses to justify funding. Conversely, suppose the bank’s threat event frequency rises, as it might if the sector experienced an elevated threat environment. Because frequency is the dominant driver of variance, this increase moves the score substantially, and the sensitivity ranking the model reports would flag frequency as the input most responsible for the change.

The same walkthrough applied to the digital trade platform tells a different story, and the difference is instructive. For the platform, an improvement in control strength reduces the direct-loss term, but because that term is a smaller share of the platform’s score, the reduction in the total is more modest than it was for the bank. What moves the platform’s score most is a change in its network centrality, because centrality

drives the trade-disruption term that dominates the platform's exposure. The lesson is that the highest-value risk reduction differs by enterprise: for the bank, strengthening internal controls yields the greatest financial return, while for the platform, reducing centrality or accelerating the containment of propagation yields more. A model that produces a single score without this sensitivity structure would obscure this lesson; the model's reporting of the sensitivity ranking is what makes it actionable, directing each enterprise's investment toward the level that moves its own exposure most.

4.5The Component Composition as a Diagnostic

The worked examples reveal a use of the model that extends beyond the headline score: the composition of the score across its components itself is a diagnostic, telling the practitioner where an enterprise's exposure concentrates and therefore where its defenses should concentrate. An enterprise whose score is dominated by the direct-loss term, as the commercial bank's is, should direct its primary effort toward strengthening internal controls that reduce the frequency and severity of its own breaches. An enterprise whose score is dominated by the trade-disruption term, as the digital trade platform's is, should direct its effort toward reducing its centrality, diversifying its dependencies, and accelerating the containment of propagation, because for such an enterprise, the largest exposure lies in the network rather than in its own systems. An enterprise whose score carries a large regulatory term should direct effort toward the controls and disclosures that mitigate regulatory consequences, including the timeliness of notification and the demonstrable quality of pre-incident controls that regulators weigh in setting penalties. The composition thus converts the score from a single number into a map of exposure, and the practitioner who reads that map can allocate defensive effort with a precision that a single number could never support. This diagnostic use is among the model's quieter virtues, and it is available only because the model preserves the separate identity of its components rather than collapsing them into an undifferentiated total.

5. Model Validation

A model that cannot be tested against reality is a hypothesis, not a tool. This section validates the Financial-Based Cybersecurity Risk Exposure Model against realized United States enterprise losses and benchmarks its accuracy against the established alternatives. The results reported here are the empirical findings of the validation exercise and are presented exactly as obtained.

5.1Validation Dataset and Enterprise Archetypes

The model was validated against one thousand two hundred and forty-seven United States enterprise cyber incidents drawn from a cyber loss database, supplemented with regulatory filings and insurance loss data covering the years 2015 to 2024. Quality filters excluded incidents with incomplete financial disclosure, unrealized losses within thirty-six months, and ambiguous sector classification. Three synthetic but empirically grounded enterprise archetypes were constructed for archetype-level validation, corresponding to the commercial bank, the cross-border digital trade platform, and the financial technology hybrid operator described in the worked examples above. These archetypes were chosen to span the range of exposure profiles the model is designed to handle; from an enterprise whose risk is concentrated in its own systems to one whose risk is dominated by propagation across a trade network.

5.2Component Contribution Analysis

The analysis of how each component contributes to the total score across the three archetypes yields the most significant single finding of the validation exercise. The contribution of the trade disruption term ranges from 18.3 percent for the commercial bank to 41.7 percent for the digital trade platform. This range is itself the result: it demonstrates that the importance of supply-chain propagation is not uniform but depends heavily on the enterprise's position in the trade network, and that for the most trade-intensive enterprises, a model that ignores propagation will understate exposure by nearly half. The regulatory cost term is consistently material across all three archetypes, never falling below roughly fifteen percent of the total, confirming that regulatory exposure is a structural feature of cyber risk in the United States rather than an occasional add-on.

Figure 5. Component contribution to the total score by enterprise archetype. Direct loss dominates for the commercial bank; trade disruption constitutes the largest single component for the digital trade platform; regulatory cost is consistently material across all archetypes. [Figure to be inserted from the original manuscript.]

5.3Monte Carlo Distribution Results

The probability distributions generated by the simulation for each archetype display the heavy-tailed character that distinguishes cyber risk from normally distributed risks. The ratio of the ninety-ninth percentile to the median averages 4.8 across the archetypes, which means that the severe-but-plausible annual loss is nearly five times the typical annual loss. This single statistic carries a profound implication for capital allocation: a risk management strategy focused on expected or typical values will systematically under-allocate capital for the tail events that actually threaten solvency. The fat tail is not a statistical curiosity; it is the part of the distribution that determines whether an enterprise survives a bad year.

Figure 6. Probability distributions of the score from the Monte Carlo simulation across ten thousand iterations for each archetype, with markers for the median, the ninetieth percentile, and the ninety-ninth percentile. The heavy-tailed structure confirms the distributional assumption. [Figure to be inserted from the original manuscript.]

5.4Accuracy Benchmarking

The accuracy of the model was compared against two established benchmarks, a benchmark based on the Factor Analysis of Information Risk model and an operational risk value-at-risk benchmark, measured by mean absolute error against realized losses. The model achieves a combined mean absolute error of 8.3 percent, compared with 30.4 percent for the Factor Analysis of Information Risk benchmark and 19.1 percent for the operational risk value-at-risk benchmark. The improvement is most pronounced for the digital trade platform, where the modelling of trade propagation captures losses that the benchmarks miss entirely. The full validation statistics are presented in the table below.

Metric	Arch. A	Arch. B	Arch. C	Combined	FAIR
Mean absolute error	6.8%	9.2%	8.9%	8.3%	30.4%
Root mean square error (thousands)	$312K	$840K	$594K	$582K	$2.1M

Direct-loss contribution	59.6%	43.5%	53.0%	52.0%	~85%
Trade-disruption contribution	18.3%	41.7%	27.4%	29.1%	None
Regulatory contribution	22.1%	14.8%	19.6%	18.9%	Partial
Ratio of ninety-ninth to median	4.2	5.6	4.7	4.8	N/A

Table 5. Validation results by archetype against the Factor Analysis of Information Risk benchmark. The model achieves a combined mean absolute error of 8.3 percent, substantially better than both benchmarks.

The most significant validation finding bears repeating, because it is the empirical heart of the paper. For the digital trade platform, the trade disruption component accounts for 41.7 percent of the total score. This confirms that ignoring supply-chain propagation for trade-intensive enterprises systematically understates financial cyber exposure by nearly half. The implication is direct and consequential: cyber insurance coverage limits, regulatory capital allocations, and control investment budgets for digital trade platforms that rely on single-entity breach cost models are structurally insufficient, leaving these enterprises and the networks they anchor materially under-protected.

5.5Reading Real Incidents Through the Model

The validation in the preceding subsections establishes that the model fits the historical record in aggregate. It is equally instructive to read a small number of well-documented real incidents through the lens of the model, asking in each case which component of the score would have dominated and what the model would have illuminated that a single-entity, qualitative assessment would have missed. The analyses that follow are interpretive rather than quantitative: they do not assign a specific score to any named organization, since doing so responsibly would require access to that organization's internal data, but they show how the structure of the model maps onto the structure of real loss events. The point is to demonstrate that the model's components correspond to real and distinct sources of financial harm, and that the model's emphasis on propagation and regulatory cost is borne out by the most consequential incidents of recent years.

5.5.1 A Health-Care Payments Processor

When a major health-care payments processor suffered a breach in 2024, the direct cost of remediation was only part of the story. The incident disrupted payment processing for a vast number of providers, propagating financial harm far beyond the breached entity into the networks that depended on it. Read through the model. This is a case in which the trade disruption term, rather than the direct-loss term, would have captured the bulk of the exposure. The breached entity sat at a position of high network centrality within the health-care payments network, and its compromise propagated to dependent organizations with a velocity that a single-entity model would entirely ignore. The model's insistence on treating propagation as a first-class component is precisely what would have allowed an analyst to anticipate that the systemic consequences of a breach at such an entity would dwarf its direct costs. The regulatory term would also have been material, given the health-privacy obligations attached to the data involved.

5.5.2 A Software Supply-Chain Espionage Campaign

The compromise of a widely used network management product in 2020, attributed to a nation-state, propagated a malicious update to thousands of organizations through a single trusted channel. This incident illustrates two features of the model at once. First, it is a supply-chain propagation event of the kind the trade disruption component is designed to capture, in which the harm spreads through a network rather than remaining contained at the point of compromise. Second, it is a nation-state event of the kind the core model addresses only through its separate fifth term, because the objective was espionage rather than financial gain, and the loss structure accordingly differs from the cybercriminal pattern. Read through the model; this incident sits at the intersection of the trade disruption component and the nation-state extension, and it illustrates why both are necessary. A model that captured only direct, financially motivated loss would have been blind to the two features that made this incident historically significant.

5.5.3 A Pipeline Ransomware Incident

When a ransomware incident disabled a major fuel pipeline operator in 2021, the operator paid a ransom, but the far larger consequence was the disruption of fuel supply across a wide region, with economic costs that exceeded the direct cost of the breach by a substantial multiple. This is a paradigmatic case for the sector multiplier and the trade disruption component working together. The energy sector carries one of the highest multipliers in the model precisely because of grid and supply interdependency, and the disruption propagated through the logistics network in a manner that the trade disruption component is built to represent. The incident demonstrates that for critical infrastructure operators, the consequences of a cyber event are governed less by the cost of restoring the breached system than by the cost of the disruption that the breach causes downstream, and the model is structured to make that downstream cost visible.

5.5.4 The Destructive Malware of 2017

The destructive malware that spread globally in 2017, generating losses estimated in the billions and prompting years of insurance litigation, is the incident that most sharply illustrates the inadequacy of questionnaire-based risk assessment and the value of distinguishing classes of threat. Insurers who had priced policies on qualitative models faced catastrophic and correlated claims they had not anticipated, and the subsequent litigation over exclusion clauses turned on whether the event was criminal or state-sponsored, a distinction the prevailing underwriting models could not represent. Read through the model, this incident vindicates two design choices: the heavy-tailed severity distribution, which assigns meaningful probability to exactly this kind of extreme correlated loss, and the threat-actor differentiation embedded in the nation-state extension, which is designed to support precisely the criminal-versus-state distinction on which the litigation turned. An insurer using the model would have been better placed both to price the tail and to anticipate the coverage question.

5.5.5 A Major Credit Bureau Breach

The breach of a major consumer credit bureau in 2017 exposed the sensitive financial data of a very large number of individuals and resulted in substantial regulatory settlements and mandated remediation. This incident is the clearest illustration of the regulatory cost component. The direct technical remediation was significant, but the defining financial consequence was the regulatory response: civil monetary penalties across multiple authorities, mandated investment in security and monitoring, and years of compliance obligation. Read through the model, the regulatory cost component would have dominated the exposure profile of this incident, and the decision-tree approach to estimating that component, mapping the categories of data, the volume of records, and the regulatory regimes involved to an expected penalty and remediation cost, would have produced an estimate of the regulatory exposure that a model omitting regulatory cost would have missed entirely.

Incident type	Dominant model component	What a single-entity model misses
Health-care payments processor	Trade disruption (high centrality)	Propagation to dependent providers
Software supply-chain espionage	Trade disruption plus nation-state term	Propagation and strategic-actor loss structure
Pipeline ransomware	Sector multiplier plus trade disruption	Downstream supply disruption
Destructive malware, 2017	Heavy tail plus nation-state differentiation	Correlated catastrophic and threat-class
Credit bureau breach	Regulatory cost component	Penalties and mandated remediation

Table 10. Five real incidents read through the model, showing the dominant component in each and what a single-entity, qualitative assessment would have missed.

Taken together, these five readings make a single point with cumulative force. The components of the model are not abstractions. Each corresponds to a real and distinct source of financial harm that has materialized in a consequential incident, and the incidents in which the largest losses occurred are precisely those in which propagation, regulatory cost, and threat-actor class, the dimensions that existing frameworks neglect, dominated the outcome. The model is, in this sense, a generalization of the lessons of these incidents: it builds into its structure the sources of harm that the most damaging events have repeatedly revealed.

5.10 Robustness and the Interpretation of Error

A combined mean absolute error of 8.3 percent invites the question of what that figure means and how much confidence it warrants. The error is measured against realized losses, which means it captures how closely the model's estimates track the actual financial outcomes of the incidents in the validation set. An error of this magnitude indicates that the model's estimates are, on average, within roughly one-twelfth of the realized loss, a level of accuracy that compares favorably with established financial risk models and that is far superior to the benchmarks, whose errors exceed nineteen and thirty percent, respectively. The error is not uniform across archetypes; it is lowest for the commercial bank, whose exposure is concentrated in well-documented direct losses, and somewhat higher for the trade platform and the financial technology hybrid, whose exposure involves the harder-to-estimate propagation and multi-sector effects. This pattern is itself reassuring, because it shows the error concentrating where the modelling is intrinsically more difficult rather than appearing randomly.

Two cautions temper the interpretation of the error. The first is that it is a retrospective measure, established by fitting historical data and comparing against known outcomes, and retrospective fit does not guarantee prospective accuracy, as the limitations section emphasizes. The second is that the error is an average, and the heavy tail of the loss distribution means that the model's estimates for the most extreme events carry greater uncertainty than the average error suggests. The confidence intervals reported alongside each estimate are designed to make this uncertainty visible, and a practitioner should attend to the width of the interval as well as to the point estimate, particularly when the decision at hand turns on the tail. An error figure, however favorable, is a summary, and the discipline of reporting the full distribution and its uncertainty is what keeps that summary from becoming a false comfort.

5.11 External Validity and Generalization

A natural question about any validation is how far its results generalize beyond the data on which they were established. The validation here was conducted against United States enterprise losses across three archetypes spanning a traditional bank, a digital trade platform, and a financial technology hybrid, and the strong fit across these archetypes provides evidence that the model generalizes across enterprises of substantially different exposure profiles. Several considerations bear on the limits of that generalization. The model is calibrated to United States regulatory and market data, and its regulatory cost component reflects the United States enforcement landscape; applying it in another jurisdiction would require recalibrating that component to the local regime, though the structure of the model would carry over. The archetypes, while empirically grounded, are synthetic constructions designed to span the relevant range, and validation against a larger and more varied set of real enterprises would strengthen the evidence further. The model is calibrated to the historical experience of financially motivated cybercrime, and its generalization to the strategic threat is handled through the separate nation-state term rather than through the core validation. Within these limits, the consistency of the fit across very different archetypes is reassuring evidence that the model captures something general about the financial structure of cyber risk rather than fitting the idiosyncrasies of a particular case.

5.12 The Choice of Archetypes Justified

The selection of the three archetypes, the commercial bank, the digital trade platform, and the financial technology hybrid, was deliberate, and the rationale for the choice bears on the strength of the validation. The three were chosen to span the range of exposure profiles the model is designed to handle, from an enterprise whose risk is concentrated in its own direct losses and regulatory obligations to one whose risk is dominated by propagation through a trade network, with a hybrid case between them. This spanning is what makes the consistent fit across the three archetypes meaningful: if the model fit well only for enterprises of a single profile, its accuracy would be of limited generality, but the fact that it fits well across profiles as different as a traditional bank and a pure trade platform is evidence that it captures something general about the financial structure of cyber risk. The archetypes are synthetic, constructed to be empirically grounded representatives of their classes rather than specific real enterprises, a choice that allows the validation to test the model against well-characterized profiles without the confounding idiosyncrasies of any single firm. A natural extension, identified among the future research directions, is to validate the model against a larger and more varied set of real enterprises, which would complement the archetype validation with evidence drawn from the full diversity of the enterprise population. The archetype

validation establishes that the model works across the principal exposure profiles; the proposed extension would establish how well it works across the full variety of real firms.

6. Extending the Model to Nation-State Threats

The model, as specified so far, calibrates its frequency and severity distributions against data on financially motivated cybercrime. This is the right foundation for the great majority of enterprises, because financially motivated incidents constitute the bulk of the loss experience. But it leaves a gap. Nation-state threat actors pursue different objectives, such as the theft of intellectual property, the disruption of critical infrastructure, the acquisition of strategic data, and their attacks generate loss structures that do not follow the cybercriminal pattern. Earlier treatments of the model acknowledged this gap and deferred it to future work. This section closes it, converting the deferral into a formally specified additional term.

6.1Why Nation-State Threats Require Separate Treatment

Nation-state cyber activity is empirically characterized by two statistical properties that distinguish it sharply from cybercriminal activity. The first is significantly lower frequency: nation-state actors mount fewer incidents, but they direct them with concentrated precision at high-value targets. The second is significantly higher severity per incident: the theft of strategic intellectual property or the disruption of critical infrastructure can generate losses that dwarf the cost of a financially motivated breach. A destructive campaign of state origin has been estimated to have caused global losses in the billions, and a major espionage campaign attributed to a nation-state has been associated with estimated economic losses an order of magnitude larger still. These are not the losses of the typical year; they are the rare, severe events that a frequency-and-severity model must represent through a separate distribution if it is to represent them at all.

Folding these events into the cybercriminal distributions would distort both. It would inflate the apparent frequency of catastrophic loss for the typical enterprise while understating the true severity of the rare state-sponsored event. The cleaner approach, and the one adopted here, is to model nation-state risk as a distinct compound process and to add it to the core equation as a fifth term, applied only to enterprises whose profile warrants it.

6.2A Fifth Term for Nation-State Exposure

The nation-state extension takes the following form, parallel in structure to the direct-loss term of the core equation but parameterized separately.

$$\text{CRES (nation-state)} = \text{TEF (ns)} \times \text{SEM (ns)} \times \text{LM (ns)}$$

The nation-state threat event frequency is parameterized from a Poisson process whose rate is calibrated from open-intelligence sources. Three such sources are feasible and complementary: the catalog of known exploited vulnerabilities maintained by the national cybersecurity agency, which documents state-attributed exploitation by sector; the annual threat reports published by major incident-response firms, which provide sector-stratified data on dwell time and attribution for state-sponsored intrusions; and publicly available trackers of state-attributed cyber operations maintained by policy research institutions. For organizations holding security clearances, classified threat intelligence from the relevant national integration centers or

sector information-sharing bodies would provide superior frequency calibration unavailable to open research.

The nation-state loss magnitude is parameterized not from empirical loss data, which is sparse and often classified, but through structured expert elicitation following an established calibrated-assessment methodology. Because expert elicitation is known to underestimate tail severity, the loss magnitude distribution for nation-state events should be specified with a heavier tail than the log-normal used for cybercriminal losses, reflecting the documented capacity of state-sponsored events to generate truly extreme outcomes. The nation-state sector multiplier applies an elevated value for the sectors most exposed to strategic targeting, particularly the defense industrial base, reflecting the severity of penalties for the loss of controlled information and the risk of contract termination, and the energy sector, reflecting the cost of restoring disrupted grid infrastructure and the associated regulatory consequences.

6.3Calibration Pathway and Scope

Full calibration of the nation-state term is identified as the highest-priority methodological extension of the model and the subject of a planned follow-on study. The pathway is clear. Frequency would be calibrated from the open-intelligence sources named above, cross-checked where possible against classified intelligence for cleared organizations. Severity would be elicited from a panel of cleared experts using a calibrated assessment protocol designed to mitigate overconfidence. The sector multipliers would be extended to reflect the distinctive consequences of strategic targeting. The follow-on study would draw on engineering frameworks developed for the protection of critical infrastructure to ground the severity estimates in the physical and operational realities of the most exposed sectors.

> **Practical guidance for affected enterprises**
>
> Enterprises in the defence, energy, and government-adjacent sectors with significant nation-state exposure should supplement the core model with the nation-state term and, where they hold clearances, with classified threat intelligence. For the great majority of commercial enterprises, the core four-component model remains the appropriate tool, and the nation-state term can be set aside. The model is modular by design: the fifth term is added only where the threat profile justifies it.

6.4An Illustrative Application of the Nation-State Term

To illustrate how the nation-state term operates, consider a defense contractor that holds controlled technical information and is therefore an attractive target for strategic espionage. For this enterprise, the core four-component model would estimate the exposure arising from financially motivated cybercrime, but it would understate the total exposure because it would not capture the strategic threat. The term " nation-state " adds a separate estimate. The nation-state threat frequency, calibrated from open-intelligence sources documenting state-attributed activity against the defense sector, would be low in absolute terms, reflecting the rarity of such events. The nation-state loss magnitude, elicited from cleared experts and specified with a heavy tail, would be large, reflecting the catastrophic potential of the theft of strategic information or the termination of a major contract. The nation-state sector multiplier for the defense industrial base would be elevated, reflecting the severity of penalties for the loss of controlled information. The product of these three would add a term to the enterprise’s score that captures the strategic exposure the core model omits. For

most commercial enterprises, this term would be negligible or absent; for the defense contractor, it could be the largest term of all, and its omission would leave the enterprise's most consequential exposure unmeasured. This illustration shows why the nation-state term is modular: it is added only where the threat profile warrants it, but where it is warranted, it can dominate.

6.5Why a Heavier Tail for Nation-State Severity

The specification of the nation-state loss magnitude with a heavier tail than the cybercriminal severity distribution is a deliberate modelling choice that deserves justification, because it materially affects the estimate of the most extreme outcomes. Cybercriminal losses, while heavy-tailed, are bounded in practice by the economic logic of the criminal: an attacker seeking financial gain has limited incentive to destroy what serves extraction, and the resulting losses, though they can be large, follow the log-normal pattern the core model uses. Nation-state losses follow a different logic. An actor pursuing strategic objectives may seek to cause maximal disruption to critical infrastructure or to extract information whose loss is catastrophic and unbounded by any extraction calculus, and the historical record contains state-associated events whose losses dwarf even the largest cybercriminal incidents. A distribution calibrated to cybercriminal losses would assign negligible probability to such events, and a model using it would be blind to exactly the outcomes that make nation-state threats consequential. Specifying the nation-state severity with a heavier tail, of the kind that assigns meaningful probability to truly extreme losses, corrects this blindness, ensuring that an enterprise exposed to strategic targeting sees in its estimate the catastrophic potential that its threat profile entails. The choice is conservative by design, because in the domain of strategic threat, the cost of understating the tail far exceeds the cost of overstating it.

6.6Integrating the Nation-State Term with the Core Model

A practical question for an enterprise that warrants the nation-state term is how to integrate it with the core four-component model without double-counting or confusion, and the integration is straightforward when approached carefully. The core model estimates exposure to financially motivated cybercrime, calibrated against the cybercriminal loss experience, and the nation-state term estimates exposure to strategic threat actors, calibrated separately. Because the two are calibrated against distinct threat populations, they can be added without double counting, much as the components within the core equation are added. The enterprise computes its core score in the usual way, computes the nation-state term where its profile warrants, and reports both the core score and the combined score, so that the decision-maker can see the exposure attributable to each class of threat. This separation is informative as well as methodologically sound, because the appropriate responses to the two classes of threat differ: defense against financially motivated cybercrime emphasizes broad control hygiene, while defense against strategic threats emphasizes the protection of specific high-value assets and, for cleared organizations, the use of classified intelligence. Reporting the two exposures separately allows the enterprise to direct its defenses appropriately rather than blurring two distinct problems into a single number, and it preserves the modularity that allows the nation-state term to be omitted entirely for the great majority of enterprises whose profile does not warrant it.

7. Sensitivity Analysis and Model Limitations

A quantitative model earns trust not by claiming certainty but by being honest about where its outputs are most uncertain and where its assumptions are most fragile. This section does both. It first decomposes the

variance of the model's output to identify which inputs matter most, and then sets out the model's limitations candidly, so that practitioners use it with appropriate caution.

7.1Global Sensitivity Analysis

A global sensitivity analysis using a variance-decomposition method was conducted across fifty thousand Monte Carlo iterations, decomposing the total variance of the score into the contributions of each model input. Three conclusions follow. First, the threat event frequency is the dominant driver of variance, accounting for roughly thirty-five percent of the total across the archetypes. This indicates that organizations should concentrate their primary data-quality efforts on accurate tracking of threat frequency through security monitoring systems and participation in industry intelligence-sharing programs. Second, the network centrality score is materially more influential for the digital trade platform than for the commercial bank, confirming that the mapping of supply-chain networks is a disproportionately high-value investment for trade-intensive enterprises. Third, the regulatory penalty parameters contribute relatively modestly, which indicates that the decision-tree approach to estimating regulatory cost provides adequate accuracy without the need for expensive multi-scenario regulatory modelling.

Figure 8. First-order sensitivity indices showing the proportion of total variance in the score attributable to each input, by archetype. Threat event frequency dominates across all archetypes; network centrality is markedly more influential for the trade platform. [Figure to be inserted from the original manuscript.]

Input variable	Arch. A	Arch. B	Arch. C	Combined
Threat event frequency	0.38	0.31	0.35	0.35
Loss magnitude distribution	0.27	0.18	0.24	0.23
Network centrality score	0.09	0.24	0.17	0.17
Sector exposure multiplier	0.14	0.08	0.11	0.11
Digital trade dependency	0.05	0.13	0.08	0.09
Regulatory penalty parameters	0.07	0.06	0.05	0.06

Table 6. First-order sensitivity indices. Values represent the proportion of total variance in the score attributable to each input independently; residual variance attributable to interactions is not shown.

7.2Model Limitations

No model is without limitations, and the credibility of a model rests in part on the candor with which its limitations are acknowledged. Five deserves particular emphasis.

7.2.1 The Treatment of Nation-State Threats

As discussed in Section 6, the core model calibrates against cybercriminal data and addresses nation-state threats through a separate term whose full calibration remains the subject of future

work. Enterprises with significant exposure to strategic targeting should treat the core model's output as a lower bound for that category of risk and supplement it accordingly.

7.2.2 The Assumption of Stationarity

The model estimates its distributional parameters from historical data and treats them as stationary over the assessment horizon. The rapid evolution of offensive capabilities augmented by artificial intelligence, including machine-assisted social engineering, automated discovery of vulnerabilities, and synthetic-media-enabled deception, may render historical frequency and severity distributions poor predictors of near-term losses. The model should therefore be recalibrated at least annually, with forward-looking adjustments to the threat frequency applied when warranted by assessments of adversary capability. Section 11 develops this point at greater length.

7.2.3 Dependence on Commercial Data

Several key inputs depend on commercial databases that carry significant annual licensing costs, which may limit accessibility for smaller enterprises, non-profit organizations, and public-sector entities operating under constrained budgets. Open-source alternatives provide partial substitutes with reduced granularity. The reduced-data variant specified in Section 10 is the model's response to this limitation, extending financial cyber risk quantification to organizations that cannot afford the full commercial data stack.

7.2.4 The Complexity of Multi-Sector Enterprises

Enterprises operating across several critical infrastructure sectors, a common situation for financial technology firms, cloud providers, and digital trade platforms, require a revenue-weighted averaging of the sector multipliers. The sensitivity analysis confirms that this weighting materially affects the score, so a single-sector approximation for a hybrid enterprise introduces systematic bias that practitioners must address through explicit multi-sector parameterization rather than convenient simplification.

7.2.5 The Limits of Retrospective Validation

The validation reported in Section 5 is retrospective: parameters were estimated from historical data and compared against known realized losses. Retrospective validation establishes that the model fits the historical record well, but it does not by itself establish predictive validity. Prospective validation, in which estimates are made before incidents occur and compared as losses materialize, would provide stronger evidence and is identified in Section 12 as the highest-priority future research initiative. Practitioners should understand the current evidence as a strong retrospective fit rather than a proven prediction.

7.3Comparative Analysis Against Existing Frameworks

It is worth setting the model explicitly against the principal existing frameworks, not to disparage them, for each was built for a purpose it serves well, but to clarify what the model adds. The comparison is best organized around the question each framework was designed to answer, because the frameworks differ less in quality than in purpose, and the model's contribution is to answer a financial question that the others were not built to address.

The Cybersecurity Framework published by the national standards body answers the question of how mature an organization's cybersecurity practices are, organizing activities into functions and describing maturity within each. It is an excellent operational guide, and the model does not replace it; rather, the model consumes its output, treating the maturity it measures as an input to the control strength factor. The information security management standards answer the question of how an organization should systematically treat its information security risks, providing a lifecycle of identification, assessment, and treatment. Again, the model complements rather than replaces them, using their risk identification as a starting point and adding the financial dimension they do not provide. The governance framework for information and related technologies answers the question of how an organization should govern its technology, operating at the level of process and capability. The model operates at a different level, that of financial quantification, and the two are compatible.

The Factor Analysis of Information Risk model is the closest existing relative, because it shares the ambition of quantifying risk in financial terms. The model in this paper extends it in three specific ways: it anchors its distributions to empirical loss data rather than relying primarily on expert elicitation, it incorporates sector-calibrated multipliers and regulatory cost components that the earlier model does not, and it treats supply-chain propagation as a first-class component. The operational risk value-at-risk models used in banking answer the question of how much capital an institution should hold against operational losses, including cyber losses, but they were not designed for the specific statistical character of cyber risk, and they do not capture propagation or regulatory cost in the way the model does. The comparison, in short, is not between a good model and bad ones, but between models built for operational, governance, and general-quantification purposes and a model built specifically to answer the financial question that securities disclosure, bank capital, operational resilience, and insurance pricing now insistently pose.

Framework	Question it answers	Relationship to the model
National cybersecurity framework	How mature are our practices?	Feeds the control strength factor
Information security standards	How do we treat our risks systematically?	Provides risk identification as a starting point
Governance framework	How do we govern our technology?	Operates at a complementary level
Factor Analysis of Information Risk	What is the risk in financial terms?	Closest relative; the model extends it
Operational risk value-at-risk	How much capital for operational loss?	Compatible; the model sharpens the cyber input

Table 12. The model in relation to the principal existing frameworks. Each framework answers a different question; the model answers the financial question the others were not built to address.

7.4Stress and Reverse Stress Testing

Beyond the sensitivity analysis, which decomposes the variance of the score under normal conditions, the model supports two forms of stress testing that probe its behavior under severe conditions, and both are valuable for the supervisory and capital applications. Forward stress testing sets the model's inputs to severe but plausible levels, for example by raising the threat frequency to a high percentile of its historical range and the loss severity to a stressed level, and reads off the resulting score, which represents the loss the enterprise might face under a severe scenario. This is the form of stress testing that supervisory programs

require, and the model's simulation architecture produces it directly. Reverse stress testing inverts the question: rather than asking what loss a given scenario produces, it asks what scenario would produce a loss large enough to threaten the enterprise's viability and works backward to the combination of inputs that would generate that loss. Reverse stress testing is valuable because it identifies the specific circumstances, the combination of elevated frequency, severe loss, and widespread propagation, that the enterprise should most fear and against which it should concentrate its defenses. The model supports reverse stress testing by allowing the analyst to fix a threatening loss level and search for the input combinations consistent with it, and the resulting scenarios often reveal dependencies, such as a reliance on a single high-centrality trade partner, that the forward analysis leaves implicit. Together, the two forms of stress testing give the enterprise a fuller understanding of its exposure to severe events than either the point estimate or the sensitivity analysis provides alone.

8. Practical Applications

The value of a financially denominated cyber risk model lies in what enables decision-makers to do. This section develops four domains of application in which the model addresses a concrete and pressing need: enterprise risk management, cyber insurance, regulatory capital and stress testing, and digital trade policy. In each, the common thread expresses cyber risk in money rather than in adjectives, unlocking decisions that were previously unavailable.

8.1Enterprise Risk Management

For the enterprise risk practitioner, the model addresses the most persistent operational limitation of the field: the inability to present cybersecurity risk in the same financial terms as credit, market, and liquidity risk. Because the output is denominated in dollars, it integrates directly into the risk registers, risk appetite frameworks, and board-level reporting that already express other categories of risk in money. This comparability enables three decisions that qualitative ratings cannot support.

The first is proportionate capital allocation. By comparing the score against the enterprise's credit and market risk estimates, a board can allocate security investment in proportion to the financial risk profile rather than relying on peer benchmarking or intuition. The second is the quantification of the return on control investment. By re-running the model with and without a proposed control, a practitioner can compute the reduction in the score per dollar of investment, producing a financially grounded prioritization metric that replaces subjective risk-reduction scoring. The third is the calibration of risk appetite. A threshold expressed as a fraction of revenue, for example a stipulation that the ninetieth-percentile score should not exceed a stated percentage of annual revenue, provides a financially meaningful and auditable boundary that a qualitative rating cannot. The securities disclosure rules add a further application: the score at the median and ninetieth percentile provides the quantitative foundation for the description of risk management processes required in annual filings, and the ninety-ninth-percentile estimate provides the financial basis for assessing the materiality of a reported incident.

8.2Cyber Insurance: Underwriting and Pricing

The cyber insurance market, with gross written premiums measured in the tens of billions of dollars, continues to rely predominantly on questionnaire-based underwriting that correlates only weakly with

realized loss outcomes. The resulting adverse selection and unexpected catastrophic losses threaten the sustainability of the market and have driven some insurers to withdraw from certain segments. The model offers three enhancements to underwriting practice. For primary underwriting, an entity-specific score replaces industry-cohort approximations, enabling premium rates that reflect the actual risk profile rather than the sector average. For the management of systemic accumulation, the trade disruption component provides a quantitative basis for understanding correlated losses across a portfolio, which is critical for insurers exposed to several high-centrality digital infrastructure enterprises whose simultaneous compromise could generate correlated claims. For the determination of coverage limits, the ninetieth- and ninety-ninth-percentile estimates anchor primary and excess limit recommendations, replacing broker benchmarking with entity-specific financial analysis. The litigation that has established the distinction between criminal and state-sponsored events as a material underwriting question is addressed directly by the threat-actor differentiation built into the model, including the nation-state extension developed in Section 6.

8.3Regulatory Capital and Stress Testing

For financial institutions subject to international bank capital accords, the model provides a cyber-specific input to the operational risk capital calculation. The regulatory cost component maps to the regulatory penalty and remediation categories of that calculation, while the direct-loss component, parameterized to the financial services multiplier, generates loss estimates consistent with the historical-loss requirements of the standardized approach. For the stress-testing programs administered by the central bank, the model's Monte Carlo architecture is directly compatible: by setting the threat frequency and loss severity to stressed levels corresponding to a severe but plausible scenario, the model generates a stress loss estimate in the format the supervisor requires. The trade disruption component is specifically relevant to supervisory concern about contagion from the compromise of a systemically important institution, a scenario the model represents explicitly through its centrality and exposure parameters. For institutions subject to the European operational resilience regime, the model's financial impact directly satisfies the requirement for documented financial consequence analysis, allowing a multinational institution to produce regulatory-grade outputs for several frameworks from a single parameterization.

8.4Digital Trade Policy

At the level of policy, the trade disruption framework enables the quantification of systemic cyber risk across national digital trade infrastructure. By aggregating scores across the enterprises that comprise a national digital trade network, weighted by their exposure and centrality, policymakers can construct an index that quantifies the total financial cyber risk embedded in that infrastructure. Such an index would support three specific applications: the identification of systemically critical digital trade enterprises whose compromise would generate disproportionate disruption and which therefore warrant targeted resilience requirements; the quantification of the financial benefit of mutual recognition agreements on cybersecurity standards, which reduce propagation velocity by accelerating cross-border incident response; and the evidence-based design of incident notification timelines within trade agreements, calibrated to the financial propagation velocities the model documents rather than to arbitrary regulatory deadlines. International trade bodies and bilateral technology councils provide institutional channels through which such analysis could inform multilateral policy.

8.5A Worked Enterprise Risk Management Decision

To make the enterprise risk management application concrete, consider a chief risk officer deciding whether to fund a proposed program of security investment, say the deployment of enhanced monitoring and identity controls, at a cost of several million dollars. Under a qualitative regime, the case for investment would rest on the assertion that it reduces risk from a higher rating to a lower one, a claim that is difficult to weigh against competing demands on the same budget. Under the model, the case becomes a calculation. The risk officer re-runs the model twice, once with the enterprise's current control strength factor and once with the improved factor the program would produce and reads off the reduction in the ninetieth-percentile score. If the program reduces the score by more than its cost, adjusted for the appropriate horizon, it is justified on financial grounds; if it does not, the budget is better spent elsewhere. The same calculation, repeated across a slate of candidate investments, produces a ranking by financial return, and the risk officer can fund down the ranking until the budget is exhausted, confident that each dollar is buying the greatest available reduction in financial risk. This is the discipline that financial denomination makes possible, and it replaces advocacy with analysis.

8.6The Logic of Insurance Pricing

The insurance application rewards a closer look at the pricing logic. An insurer pricing a cyber policy must estimate the expected loss it will bear under the policy, load that estimate for the cost of capital it must hold against the risk and add a margin. The expected loss is the integral of the loss distribution over the layer the policy covers, and the cost of capital is governed by the tail of that distribution, because it is the tail that determines how much capital the insurer must hold to remain solvent against a bad year. The model supplies both inputs directly. The ninetieth-percentile score informs the attachment and limit of the primary layer; the ninety-ninth-percentile score and the expected shortfall beyond it inform the capital load and the design of any reinsurance. Crucially, the trade disruption component allows the insurer to understand correlation across its portfolio: two enterprises that depend on the same digital trade infrastructure are not independent risks, and an insurer that prices them as independent will be surprised by a correlated loss. The model makes the correlation visible by exposing the shared centrality, and an insurer that prices for it will hold adequate capital where one that ignores it will not.

8.7Mapping the Model to Regulatory Requirements

The regulatory applications can be made precise by mapping each output of the model to the specific requirement it serves. The table below sets out that mapping for the principal regimes, and it is intended as a practical reference for a compliance officer assembling the documentation a supervisor expects.

Regulatory requirement	Model output that serves it	Use
Securities incident materiality	Ninety-ninth-percentile score	Financial basis for the materiality judgment
Securities periodic disclosure	Median and ninetieth-percentile score	Quantitative description of risk management
Bank operational risk capital	Direct-loss term plus regulatory capital adjustment	Input to the loss component of the capital calculation
Bank stress testing	Score under stress frequency and severity	Scenario loss estimates in required format

Operational resilience impact assessment	Full score with component breakdown	Documented financial consequence analysis

Table 11. Mapping the outputs of the model to the principal regulatory requirements they serve. A single parameterization produces the outputs for several regimes at once.

8.8A Deployment Scenario

It is useful to close the applications section with a brief sketch of how the model would be deployed within an enterprise over the course of a year, because the value of the model is realized through a cadence of use rather than a single calculation. At the start of the year, the risk team produces a baseline score and presents it to the board alongside the other enterprise risk metrics, establishing the financial magnitude of cyber risk for the year. During the year, the score is refreshed quarterly, with the trajectory reported to the board so that any deterioration is visible early. When a significant control investment is proposed, the model is run to quantify its financial return, informing the funding decision. When an incident occurs, the model provides the financial basis for materiality assessment and the regulatory notifications that follow. At the annual cycle's close, the model is recalibrated against the year's data, and the cycle begins again. Deployed in this way, the model is not a report produced once and filed, but a living instrument woven into the governance of the enterprise, and it is in this woven form that its value is greatest.

8.9Mergers, Acquisitions, and Due Diligence

A further application, less obvious than the regulatory and insurance uses but increasingly important, is in the due diligence that precedes a merger or acquisition. When one enterprise acquires another, it acquires the target's cyber risk along with its assets, and a target with poorly understood cyber exposure can transform a sound acquisition into a costly one. The model provides a structured way to quantify the target's cyber exposure as part of due diligence, producing a score that can be weighed against the purchase price and the projected synergies. A target whose score is high relative to its value carries a hidden liability that should be reflected in the price or addressed through remediation before completion. The model also illuminates integration risk, because connecting the acquirer's and the target's networks can increase the centrality of both and thereby raise the combined entity's trade disruption exposure above the sum of the parts. An acquirer that quantifies these effects in advance can structure the transaction and the integration to manage them, rather than discovering them after the deal has closed.

8.10 Third-Party and Vendor Risk

The same logic that makes the trade disruption component valuable for assessing an enterprise's own exposure makes it valuable for assessing the exposure that vendors and third parties introduce. An enterprise that depends on a vendor occupying a high-centrality position in a shared network inherits a portion of that vendor's risk, because the vendor's compromise would propagate to the enterprise. The model allows an enterprise to quantify this inherited risk by mapping its vendor relationships into its trade network and computing the contribution of each vendor to its own score. This quantification supports a more disciplined approach to vendor risk management than the questionnaires on which it commonly relies, allowing the enterprise to concentrate its attention and its contractual protections on the vendors whose

compromise would cost it most, rather than spreading its effort uniformly across vendors of very different systemic importance.

8.11 Educating the Board

A subtler but genuine application is the education of the board. Boards are increasingly expected to exercise oversight of cyber risk, but many board members lack the technical background to engage with cyber risk presented in technical terms, and the qualitative heat maps that have served as the lingua franca of cyber risk reporting offer them little to grasp. A financially denominated score, presented alongside the other financial risk metrics the board already understands, gives board members a handle on cyber risk that connects to their existing financial intuition. The trajectory of the score over time, the comparison of the score against the enterprise's appetite, and the financial return of proposed control investments are all things a board can engage with substantively, and the model thereby raises the quality of board oversight from the passive reception of technical briefings to the active governance of a financial risk. This educational effect is not the model's primary purpose, but it is a real and valuable consequence of expressing cyber risk in the language the board already speaks.

8.12 Implications for Practice

Stepping back from the individual applications, several implications for practice follow from the model. The first is that cyber risk can and should be brought onto the same financial footing as the other risks an enterprise manages, and that the persistence of qualitative cyber risk reporting is a choice rather than a necessity. The second is that the exposure of an enterprise cannot be understood in isolation from the network it inhabits, because for an increasing share of enterprises, the largest exposure is the propagation that a single-entity view cannot see. The third is that regulatory cost is a structural and predictable component of cyber exposure in the United States, not an occasional misfortune, and that a practice which omits it systematically understates exposure. The fourth is that the discipline of reporting a distribution and its uncertainty, rather than a single number, is what distinguishes a credible practice from a misleading one, because the tail is where solvency is decided, and the tail is precisely what a point estimate hides. A practice that internalizes these four implications will manage cyber risk more soundly than one that does not, whether it adopts this model.

8.13 Implications for Policy

At the level of policy, the model carries implications that extend beyond the individual enterprise. The capacity to aggregate exposure across the enterprises that comprise the national digital trade infrastructure offers policymakers a quantitative basis for identifying where systemic cyber risk concentrates, a basis they currently lack. This capacity could inform the designation of systemically critical digital trade enterprises, analogous to the designation of systemically important financial institutions, and the targeted resilience requirements such designation would justify. It could quantify the financial benefit of international cooperation on cybersecurity standards and incident response, giving negotiators a concrete figure to weigh against the costs of coordination. And it could ground the design of incident notification timelines in the measured velocities of financial propagation rather than in arbitrary deadlines. More broadly, the model embodies a policy proposition: that the financial measurement of cyber risk is a public good as well as a private one, because the systemic risk embedded in shared digital infrastructure is borne collectively, and

that public investment in the data and standards required for such measurement would yield returns across the economy. The talent required to perform this measurement, and the governance required to use it responsibly, are themselves matters of public concern, as the next section discusses.

8.14 Risk Transfer to the Capital Markets

Beyond traditional insurance, the financial denomination the model provides opens a path to the transfer of cyber risk to the capital markets, a development that several market participants have begun to explore as the traditional cyber insurance market strains under the weight of correlated and catastrophic losses. Instruments that transfer risk to investors, structured so that investors bear losses above a defined threshold in exchange for a yield, require exactly the kind of credible, probabilistic loss estimate the model produces, because investors will not bear a risk they cannot price. The ninety-ninth percentile estimates and the expected shortfall beyond it provide the basis for structuring the threshold and the yield, and the trade disruption component provides the basis for understanding the correlation that determines whether a portfolio of such instruments is genuinely diversified or secretly concentrated. By bringing cyber risk within reach of the deep capacity of capital markets, the financial measurement model provides could relieve the pressure on the traditional insurance market and expand the total capacity available to bear cyber risk, a systemic benefit that extends well beyond any single enterprise. This application remains nascent, and it depends on the maturation of the measurement methods this paper develops, but it illustrates how the consequences of credible financial measurement ramify through the entire architecture of risk transfer.

8.15 Placing Cyber Risk Alongside Credit and Market Risk

The deepest practical consequence of expressing cyber risk in money is that it can finally take its place alongside credit risk and market risk in the unified view of enterprise risk that boards and regulators increasingly demand. In a mature financial institution, credit risk is measured in expected and unexpected losses, market risk in value-at-risk and expected shortfall, and operational risk in capital held against historical and modelled losses. Cyber risk has stood apart from this view, measured in a different language and reported through a different channel, with the result that it cannot be aggregated with the other risks or weighed against them in the allocation of capital and attention. The model dissolves this separation. Because its output is denominated in the same units as the other risk measures and reports the same quantiles, the value-at-risk and the expected shortfall, the cyber risk estimate can be placed directly into the institution's aggregate risk view, aggregated with the other risks subject to the correlations between them, and weighed against them in the allocation of capital. This integration is not merely a matter of tidiness; it is what allows an institution to manage its total risk coherently rather than managing cyber risk in a silo divorced from the rest. The aspiration to a unified view of enterprise risk has been frustrated for cyber risk precisely by the absence of a common measure, and the provision of that measure is among the model's most consequential contributions to the practice of risk management.

8.16 The Cost of Inaction

It is worth stating directly what an enterprise forgoes by declining to measure its cyber risk in financial terms, because the cost of inaction is itself an argument for the model. An enterprise that manages cyber risk qualitatively cannot allocate its security budget by financial return, and so it will tend to overspend on visible or fashionable risks and underspend on the propagation and regulatory exposures that the headline incidents repeatedly show to dominate. It cannot price the risk it transfers to insurers, and so it will accept coverage limits set by broker benchmarking that may bear little relation to its actual exposure, leaving it underinsured against the tail or overpaying for the body of the distribution. It cannot discharge its regulatory obligations with the financial rigor those obligations increasingly demand, and so it exposes itself to enforcement for inadequate disclosure or capital planning. And it cannot give its board the financial view of cyber risk that sound governance requires, and so it leaves its board to oversee a material risk through a lens that obscures the magnitude of what is at stake. None of these costs is hypothetical; each has materialized in the incidents and enforcement actions of recent years. The cost of inaction, in short, is not merely the absence of a useful tool but the persistence of systematic errors in budgeting, insurance, compliance, and governance, errors that the financial measurement of cyber risk is precisely designed to correct. Against this cost, the effort to implement the model, however substantial, is a sound investment.

8.17 A Note on Benchmarking Against Peers

Practitioners often wish to benchmark their cyber risk against that of their peers, and the model supports such benchmarking while guarding against its misuse. Because the model produces a financially denominated score with a transparent component structure, an enterprise can compare its score and the composition of that score against the scores of comparable enterprises, identifying whether its exposure is unusually high or unusually concentrated in a particular component. This comparison can be illuminating, revealing, for example, that an enterprise carries far more trade-disruption exposure than its peers because of an unusual centrality it had not recognized. The caution is that benchmarking should inform rather than replace the enterprise's own analysis, because the peer average is not a target, and a score below the peer average does not establish that an enterprise's exposure is acceptable. The appropriate level of cyber risk for an enterprise is determined by its own risk appetite and capital, not by the behavior of its peers, and an enterprise that managed its cyber risk to the peer average would be allowing the collective behavior of its industry to set a standard that may be too lax or too stringent for its own circumstances. Used to inform the enterprise's own judgment, benchmarking is valuable; used to substitute for that judgment, it reintroduces precisely the herd behavior that the financial measurement of risk is meant to replace. The model, by grounding the score in the enterprise's own exposure, keeps benchmarking in its proper supporting role.

9. Implementation Guidance for Risk Practitioners

A model that cannot be implemented is a curiosity. This section provides practical guidance on the data the model requires, the governance processes into which its outputs should be integrated, and a maturity model that allows organizations to adopt the model at a level of precision consistent with their obligations and resources.

9.1Data Requirements and Collection

Successful implementation requires six categories of enterprise-specific data. The first is historical internal loss data for calibrating threat frequency, available from security monitoring and incident-tracking systems.

The second is an inventory of security controls and their maturity, available from assessments against established control catalogs. The third is a mapping of digital trade partners for computing network centrality, available from procurement, treasury, and supply-chain systems. The fourth is the attribution of revenue by digital channel for computing the dependency ratio, available from financial reporting. The fifth is the prior regulatory history of the enterprise, available from legal and compliance records. The sixth is the sector classification, which follows deterministically from the enterprise's business description. For organizations without comprehensive historical loss data, the model supports a hybrid approach in which enterprise-specific data is used where available and sector-cohort data is substituted where it is not. Because the sensitivity analysis identifies threat frequency as the highest-impact parameter, investment in internal loss tracking is the single highest-priority data investment for an organization beginning implementation.

9.2Governance and Reporting Integration

The output of the model should be integrated into three governance processes. At the board level, the ninetieth-percentile score and its trajectory over time should be reported quarterly alongside other enterprise risk metrics, enabling the board to discharge its oversight responsibility with financially grounded data. At the executive level, the ranking of inputs by their influence on the score should inform the annual prioritization of control investment, connecting security decisions to financial risk reduction with an auditable foundation. At the level of the audit committee, the width of the confidence interval provides a measure of model uncertainty that auditors can assess alongside the point estimate, supporting the kind of quality evaluation that rigorous committees increasingly demand. Reporting a distribution and its uncertainty, rather than a single number, is itself a discipline that strengthens governance.

9.3An Implementation Maturity Model

Organizations at different levels of risk management maturity will implement the model at different levels of precision. A three-tier maturity model guides this choice. At the foundational tier, threat frequency is drawn from sector cohort data, the loss distribution is a sector average, and the output is an order-of-magnitude estimate suitable for board communication and the setting of insurance limits. At the intermediate tier, internal monitoring data is blended with sector data, the loss distribution is adjusted to the enterprise, and the output supports disclosure language and insurance underwriting submissions. At the advanced tier, real-time threat intelligence is combined with internal data, the loss distribution is fitted and updated regularly, the network is mapped in full, and the output, accompanied by confidence intervals, is suitable for regulatory capital submissions and enforcement-grade materiality assessments. Organizations should target the tier consistent with their regulatory obligations.

Dimension	Foundational	Intermediate	Advanced
Threat frequency source	Sector cohort data	Internal monitoring blended with sector data	Real-time intelligence plus internal data
Loss distribution	Sector average	Enterprise-adjusted	Fitted, updated regularly
Output precision	Order of magnitude	Within roughly twenty-five percent	Within roughly ten percent, with intervals
Primary use	Board communication	Insurance and risk integration	Regulatory capital and stress tests

Indicative effort	Four to eight-person weeks	Twelve to twenty person-weeks	Twenty-four to forty-person-weeks

Table 7. The implementation maturity model. Each tier specifies the data sources, output precision, and primary use case appropriate to a level of risk management investment. Regulatory-grade use requires an advanced tier.

A consistent pattern emerges across implementations: the dominant expenditure of effort at the foundational tier is the extraction and cleaning of data, typically most of the total effort, rather than the development of the model itself. At the intermediate tier, the bottleneck shifts to fitting sector-specific distributions and approximating network centrality. At the advanced tier, the construction of the full network graph, while demanding, can be accelerated through commercial supply-chain intelligence platforms. Organizations adopting the model for regulatory capital purposes should additionally budget for independent model validation by a qualified third party, consistent with established supervisory guidance on model risk management. Building the data pipeline, in short, is real work; mathematics is the easy part.

9.4Building the Data Pipeline

Because the dominant expenditure of effort in implementation is the assembly of data, the construction of the data pipeline deserves specific guidance. The pipeline has five stages. The first is extraction, in which lost data is drawn from incident-tracking and monitoring systems, control data from assessment tools, trade data from procurement and treasury systems, financial data from reporting systems, and regulatory history from legal and compliance records. The second is cleaning, in which the extracted data is normalized, deduplicated, and checked for completeness, a stage that typically consumes most of the effort because enterprise data is rarely as structured as a model requires. The third is mapping, in which each data element is mapped to the model parameter it informs. The fourth is fitting, in which the distributions are estimated from the mapped data. The fifth is validation, in which the assembled parameterization is checked against any available realized losses. An organization that invests in automating this pipeline, rather than assembling the data by hand each cycle, converts a periodic project into a repeatable process, and it is the repeatability that makes the quarterly cadence described above sustainable.

9.5Governance Artifacts and Cadence

Sound governance of the model requires a small set of artifacts maintained on a regular cadence. A model documentation file records the parameterization, the data sources, and the assumptions, so that a third party can reproduce the result, a requirement of supervisory model risk management guidance. A calibration log records each recalibration and the reason for any change in parameters, establishing an audit trail. A results record records each score produced, with its confidence interval and the date of production, so that the trajectory over time is preserved. An assumptions review, conducted at least annually, revisits the central assumptions of the model, considering the evolving threat landscape and records any adjustments. These artifacts are not bureaucratic overhead; they are what allow the model to withstand the scrutiny of an auditor or a supervisor, and an organization that maintains them from the outset will find the regulatory use of the model far smoother than one that assembles them retrospectively under examination pressure.

9.6Common Pitfalls

Several pitfalls occur in implementations, and naming them helps practitioners avoid them. The first is over-precision: reporting a score to a false number of significant figures invites confidence the model does not warrant, and results should be reported with confidence intervals so that the uncertainty is visible. The second is stale calibration: a model calibrated once and left unchanged will drift away from reality as the threat landscape evolves, which is why annual recalibration is a requirement rather than a recommendation. The third is a single-sector approximation for multi-sector enterprises, which introduces systematic bias and must be corrected through explicit revenue-weighted multipliers. The fourth is neglect of the trade disruption component for enterprises that appear, superficially, to have low trade dependency but in fact sit at high centrality within a critical network, a neglect that the health-care payments case shows can be catastrophic. The fifth is the treatment of the model's output as a substitute for judgment rather than an input to it; the model informs decisions, but the decisions remain the responsibility of accountable human beings.

9.7A Ninety-Day Adoption Plan

For an organization beginning from a standing start, a structured ninety-day plan brings the model into productive use without overwhelming the team. In the first thirty days, the organization assembles the foundational data, extracts and cleans the loss, control, and financial data, and produces a first foundational-tier score, accepting that it will be an order-of-magnitude estimate. In the second thirty days, the organization refines the calibration, blends internal data with sector data, maps the principal trade dependencies, and produces an intermediate tier score suitable for board communication and insurance discussions. In the final thirty days, the organization integrates the model into governance, establishes the quarterly cadence, produces the documentation artifacts, and presents the first formal score to the board. At the end of ninety days, the organization has a working model, a sustainable process, and a baseline against which future quarters can be measured. The plan deliberately front-loads a usable result, on the principle that a rough estimate in hand at thirty days does more to build organizational confidence than a precise estimate promised at one hundred and eighty.

9.8Effort, Cost, and Team Composition

Organizations planning an implementation benefit from realistic estimates of the effort, cost, and team composition each maturity tier requires, and the estimates below are derived from implementation experience and from benchmarking against cyber risk consulting practice. They assume the organization has existing monitoring infrastructure generating structured data, a risk analyst or quantitative professional available, and executive sponsorship sufficient to secure data access and cooperation across business units.

Dimension	Foundational	Intermediate	Advanced and regulatory-grade

Team	One analyst, part-time	One full-time analyst plus a part-time data engineer	Several analysts, a data scientist, an architect, and compliance support
Effort (person-weeks)	Four to eight	Twelve to twenty	Twenty-four to sixty and beyond
Indicative external cost	Tens of thousands of dollars	Roughly one hundred to two hundred thousand	Several hundred thousand to over a million
Key bottleneck	Data extraction and cleaning	Distribution fitting and network approximation	Full network mapping and independent validation

Table 13. Indicative effort, cost, and team composition by maturity tier. The dominant expenditure at the foundational tier is data extraction and cleaning, not model development.

Two patterns recur across these estimates and deserve emphasis. The first is that the dominant cost at the lower tiers is data work rather than modelling, which means that an organization's existing data maturity is the strongest predictor of how quickly and cheaply it can adopt the model. The second is that the step from the advanced tier to a regulatory-grade implementation is driven less by the modelling itself than by the documentation and independent validation that supervisory expectations require, costs that are real but that buy the credibility a regulatory use demand. Organizations should therefore plan their adoption around their data infrastructure and their regulatory obligations rather than around mathematics, which is, as noted earlier, the least expensive part of the undertaking.

9.9Change Management and Organizational Adoption

The technical implementation of the model is necessary but not sufficient for its successful adoption, because a model that is not trusted and used by the people whose decisions it is meant to inform delivers no value, however well it is built. The organizational adoption of the model is a change management undertaking, and it succeeds or fails on factors that have little to do with mathematics. The first factor is sponsorship: the model must be championed by a leader with the authority to secure data access and to require that the model's outputs be used in decisions, because without that authority, the model will be produced and ignored. The second factor is trust: the people who must act on the model's output must understand its logic well enough to trust it, which is why the model's emphasis on auditability and on the clear explanation of its components matters as much as its accuracy. The third factor is integration: the model's output must be woven into the existing rhythms of governance, the board reporting, the budget process, and the audit cycle, rather than presented as a separate exercise, because a number that does not connect to a decision will not survive. The fourth factor is humility: the model must be presented as an input to judgment rather than a replacement for it, because a model presented as infallible will be discredited by its first error, while a model presented as a useful but imperfect instrument will survive the inevitable surprises. An organization that attends to these four factors will adopt the model successfully; one that attends only to the technical implementation will build a model that no one uses.

9.10 Tooling and Software Architecture

The implementation of the model in software deserves brief practical guidance, because the choice of tooling affects the cost, the reproducibility, and the auditability of the result. The core of the model is a Monte Carlo simulation, which can be implemented in any general-purpose analytical environment, and the simulation

itself is computationally modest, well within the capacity of a standard analytical workstation even at the fifty-thousand-iteration resolution used for sensitivity analysis. The components surrounding the simulation, the calibration of parameters, the decision tree for regulatory cost, and the network analysis for centrality are likewise standard analytical tasks. The architectural choices that matter most are not about computational power but about reproducibility and governance. The parameterization should be stored in a form that a third party can inspect and rerun, so that an auditor or supervisor can reproduce the result, a requirement of model risk management. The data pipeline should be automated so that recalibration is a repeatable process rather than a manual project, which is what makes the quarterly cadence sustainable. The results should be versioned, so that each score can be traced to the parameterization and data that produced it. And the simulation should use a recorded random seed where reproducibility of the exact figures is required, so that a result can be regenerated identically. These architectural disciplines cost little to adopt at the outset and are expensive to retrofit, so an organization building the model in software should adopt them from the first version rather than discovering their necessity under the pressure of an examination.

9.11 Maintaining the Model Over Time

A model is not a deliverable produced once but an instrument maintained over time, and the practices of maintenance determine whether the model remains trustworthy or drifts into obsolescence. The central maintenance practice is recalibration, conducted at least annually, in which the model's parameters are re-estimated against the most recent data so that they reflect the current threat landscape rather than a historical one. Recalibration is especially important in a domain evolving as quickly as cyber risk, where the offensive use of artificial intelligence and the growth of digital trade can shift the frequency and severity of losses materially within a single year, and a model calibrated to last year's experience may understate this year's risk. Alongside recalibration, the model's assumptions should be reviewed periodically, because an assumption that was sound when the model was built may be strained by developments the model's designers did not anticipate, and the conscious revisiting of assumptions guards against the silent accumulation of error. The data pipeline should be monitored for quality, because the model's accuracy depends on the integrity of its inputs, and degraded data produces degraded estimates regardless of the soundness of the model itself. And the model's performance should be tracked against realized losses as they occur, so that any systematic divergence between estimate and outcome is detected and corrected. Maintained with these practices, the model remains a living and trustworthy instrument; neglected, even the best-built model will drift away from the reality it is meant to measure, and the discipline of maintenance is therefore as essential to the model's value as the soundness of its original construction.

9.12 Independent Validation and Assurance

For any use of the model that carries regulatory, capital, or insurance weight, independent validation is not optional but essential, and the practice of independent validation deserves explicit description because it is what converts an internally credible model into one a supervisor or auditor can rely upon. Independent validation means that a party other than the team that built the model examines its design, its data, its calibration, and its outputs, and forms an independent view of whether the model is sound and fit for its purpose. The validator checks that the distributional assumptions are supported by the data, that the calibration reproduces the historical record within the stated error, that the parameters are reasonable, and that the documentation is sufficient for a third party

party to reproduce the result, and that the model's limitations are clearly stated and appropriately reflected in its use. Established supervisory guidance on model risk management sets out the expectation that material models be subject to such validation, and the model in this paper is designed to support it through the auditability of its components, the transparency of its calibration, and the discipline of its documentation. An enterprise that submits the model's outputs for a regulatory or capital purpose should budget for independent validation by a qualified party and should treat the validation not as a hurdle to be cleared but as a genuine check on the soundness of a model on which consequential decisions will rest. A model that has survived rigorous independent validation carries credibility that no internal assurance can match, and the investment in validation is repaid in the confidence with which the model's outputs can then be used.

10.A Reduced-Data Variant for Smaller Enterprises

The full model draws on commercial data subscriptions and specialized analytical capacity that smaller enterprises may lack. Yet small and medium enterprises, of which there are tens of millions in the United States, form a large and analytically underserved part of the digital trade infrastructure, and their exclusion from financial cyber risk quantification leaves both them and the networks they participate in less protected. This section specifies a reduced-data variant designed to extend the model to these organizations, trading some precision for accessibility while preserving the model's essential logic.

10.1 Design Principles

The reduced-data variant rests on three principles. First, every commercial data input is replaced with a publicly available substitute wherever a substitute exists. Threat frequency is drawn from the public sector-cohort tables of the annual breach investigations literature rather than from a commercial threat intelligence subscription. Loss distributions are fitted to public breach-portal records and a sample of regulatory filings rather than to a commercial loss database. Regulatory cost is estimated from public enforcement announcements rather than a proprietary enforcement database. Network centrality, the hardest input to approximate without commercial supply-chain mapping, is estimated from self-reported vendor lists and public procurement data and is used only at the higher tiers where trade exposure is material. Second, the variant is honest about the accuracy it sacrifices, reporting an expected inflation of error for each substitution so that users understand the precision of their estimates. Third, the variant preserves the full model's structure, so that an organization can upgrade input by input as its analytics investment grows, rather than facing an all-or-nothing choice.

10.2 Accuracy Trade-Off

Sensitivity testing, in which commercial inputs were progressively replaced with their open-source equivalents and the error re-measured against the validation dataset, indicates that the reduced-data variant achieves a combined mean absolute error in the range of roughly eighteen to twenty-six percent, compared with the 8.3 percent of the full model. This is a meaningful loss of precision, but it leaves the variant substantially more accurate than the benchmark based on the Factor Analysis of Information Risk model, whose error exceeds thirty percent, and broadly comparable to the operational risk value-at-risk benchmark. For an organization that would otherwise have no financial cyber risk estimate at all, an estimate accurate

to within roughly a fifth to a quarter is a substantial improvement over a qualitative rating, and it is sufficient for the two uses smaller organizations most need: communicating risk to leadership and setting insurance limits.

Input	Full model source	Reduced-data substitute	Error inflation
Threat frequency	Commercial threat intelligence	Public breach-investigation cohort tables	Six to nine points
Loss distribution	Commercial loss database	Public breach portal plus filings sample	Ten to sixteen points
Sector sub-scores	Commercial pricing intelligence plus full enforcement database	Public enforcement announcements plus open economic tables	Four to seven points
Regulatory cost	Full proprietary enforcement database	Public enforcement announcements	Five to eight points
Network centrality	Commercial supply-chain mapping	Self-reported vendor list plus procurement data	Eight to fourteen points
Combined error	8.3 percent (validated)	Roughly eighteen to twenty-six percent	Acceptable for foundational use

Table 8. The reduced-data variant: substitutions and the estimated inflation of error for each. The combined error remains substantially better than the Factor Analysis of Information Risk benchmark, confirming practical utility despite the accuracy cost.

The practical recommendation is straightforward. Smaller organizations should implement the reduced-data variant at the foundational or intermediate tier and upgrade to commercial data sources as their analytics investment scales. The variant is not a permanent compromise but an on ramp, designed so that the first estimate an organization produces is also the first step toward a more precise one.

10.3 A Worked Example for a Small Enterprise

To show the reduced-data variant in operation, consider a small importer with annual revenue of a few tens of millions of dollars, a modest digital trade dependency, and no access to commercial cyber risk data. Under the full model, this enterprise cannot be assessed without subscriptions it cannot afford. Under the reduced-data variant, it can. Its threat frequency is taken from the public sector-cohort tables for its industry, adjusted for its size. Its loss distribution is fitted to public breach-portal records and a sample of regulatory filings for enterprises of comparable scale. Its sector multiplier sub-scores are estimated from public enforcement announcements and open economic tables. Its regulatory cost is estimated by public enforcement records. Its network centrality, the hardest input, is approximated from its self-reported vendor list and public procurement data, and because its centrality is low, the approximation introduces little error into the final score. The variant produces a score accurate to within roughly a fifth to a quarter, which the importer uses for two purposes: to communicate to its leadership the financial magnitude of its cyber risk, and to set the

limit of the cyber insurance policy it purchases. Neither use requires the precision of the full model, and the variant's estimate is a vast improvement over the qualitative rating that would otherwise be the importer's only option. As the importer's analytics investment grows, it can upgrade input by input toward the full model, and the variant's preservation of the full model's structure makes that upgrade a matter of replacing data sources rather than rebuilding from scratch.

10.4 The Upgrade Path in Detail

The design of the reduced-data variant as an on-ramp rather than a permanent compromise rewards a closer look at the upgrade path, because the path is what allows an organization to grow its analytical capability incrementally rather than facing a single daunting leap. The path proceeds input by input, in order of the value each upgrade adds. The first upgrade an organization should pursue is to its loss distribution, replacing the public breach-portal data with enterprise-specific loss data as the organization accumulates it, because the loss distribution is among the highest impact inputs, and enterprise-specific data sharpens it considerably. The second upgrade is to threat frequency, replacing the public cohort tables with internal monitoring data blended with commercial intelligence, because frequency is the dominant driver of variance and the upgrade that most improves precision. The third upgrade, relevant only for organizations with material trade exposure, is to network centrality, replacing the self-reported vendor approximation with commercial supply-chain mapping, because for high-centrality organizations, this is the input whose approximation introduces the most error. At each step, the organization replaces one data source and leaves the rest of the model unchanged, and at each step, the error falls toward the eight-and-three-tenths-percent benchmark of the full model. The upgrade path thus converts the daunting question of whether to adopt the full model into the manageable question of which input to upgrade next, and it ensures that an organization captures value at every stage rather than only at the end.

10.5 Equity and Access in Cyber Risk Measurement

The reduced-data variant raises a consideration that is as much about fairness as about methodology: the equity of access to sound cyber risk measurement. The capacity to measure cyber risk well has, until now, been concentrated among the largest and best-resourced organizations, those able to afford the commercial data and specialized staff the full model requires, while the small and medium enterprises that make up the great bulk of the economy have been left with qualitative ratings or nothing at all. This concentration is inequitable because small enterprises bear cyber risk no less than large ones and suffer disproportionately when they bear it unknowingly, and it is inefficient, because small enterprises form an integral part of the digital trade networks whose systemic risk depends on the resilience of every node. The reduced-data variant is, in part, a response to this inequity, demonstrating that useful financial cyber risk measurement is possible from public data and within modest resources, and thereby extending the capacity for sound measurement to organizations previously excluded from it. The broader aspiration, consistent with the case for treating cyber risk data as shared infrastructure, is a future in which sound cyber risk measurement is accessible to every organization that bears the risk, not merely to those that can afford the most expensive tools. The reduced-data variant is a step toward that future, and its development reflects a conviction that the benefits of financial cyber risk measurement should be broadly shared rather than concentrated among the few.

11.Artificial Intelligence on Both Sides of the Contest

No treatment of modern cyber risk would be complete without confronting the role of artificial intelligence, which is reshaping both the offensive threat and the defensive response at a pace that strains the stationarity assumption of any historical model. This section examines artificial intelligence as a force on both sides of the contest and explains how the model accommodates it, while acknowledging the genuine uncertainty that this rapidly moving frontier introduces.

11.1 Artificial Intelligence as an Offensive Force

On the offensive side, artificial intelligence lowers the cost and raises the effectiveness of several classes of attack. Machine-assisted social engineering allows adversaries to craft persuasive and personalized deceptive messages at scale, eroding the protection that awkward phrasing once provided against fraudulent communications. Automated discovery of vulnerabilities accelerates the pace at which weaknesses are found and exploited. Synthetic media, including convincing imitations of voices and faces, enables new forms of deception in which an instruction that appears to come from a trusted executive is in fact fabricated. Each of these developments tends to increase the frequency of loss events and, in some cases, their severity, which means that historical frequency and severity distributions calibrated before the widespread availability of these capabilities may understate near-term risk. The model addresses this through its provision for forward-looking adjustment of the threat frequency parameter, applied when assessments of adversary capability warrant it, and through its requirement of at least annual recalibration.

11.2 Artificial Intelligence as a Defensive Force

On the defensive side, artificial intelligence augments the detection of threats and the automation of response, allowing security teams to identify and contain incidents faster and to relieve human analysts of routine work so that they can concentrate on complex and strategic threats. Within the financial sector specifically, institutions are increasingly approaching the governance of artificial intelligence in a manner analogous to the management of model risk, recognizing that an artificial intelligence system used for fraud detection or credit decisions is itself a model whose behavior must be validated, monitored, and controlled. This governance dimension is precisely where the expertise underlying this paper is concentrated, and it represents one of the most consequential frontiers in financial cybersecurity, because a defensive tool that is poorly governed can itself become a source of risk through bias, opacity, or adversarial manipulation.

11.3 The Governance Imperative in Financial Systems

The governance of artificial intelligence in financial systems has moved from a theoretical concern to an explicit priority of public policy. National financial authorities have issued guidance recognizing the opportunities and the risks that artificial intelligence presents to the security and resilience of the financial sector, calling on institutions to strengthen their governance of these systems, to expand the training of staff in their safe use, and to adapt established risk management frameworks to the specific demands of finance. This is a direct acknowledgement, at the level of national policy, that the responsible governance of artificial intelligence in finance is a matter of importance and that the supply of professionals capable of providing it is a recognized constraint. The model accommodates this frontier in two ways: by treating the artificial intelligence systems deployed in defense as components whose governance maturity feeds into the control

strength factor, and by recognizing, through its forward-looking frequency adjustments, that the offensive use of artificial intelligence is reshaping the threat landscape faster than historical data alone can capture. The honest position is that this is the area of greatest genuine uncertainty in the model, and the area where ongoing recalibration and expert judgment matter most.

11.4 Governing Artificial Intelligence as Model Risk

When a financial institution deploys an artificial intelligence system to detect fraud, score credit, or screen transactions, that system is, in the language of bank supervision, a model, and it falls within the established discipline of model risk management. This framing is powerful because it brings a mature governance apparatus to bear on a new technology. The discipline of model risk management requires that a model be validated independently before deployment, that its performance be monitored continuously, that its limitations be documented, and that its use be governed by clear accountability. Applied to an artificial intelligence system, these requirements translate into concrete practices: validating that a fraud-detection model performs as claimed across the populations it will encounter, monitoring for the drift that occurs as adversaries adapt, documenting the conditions under which the model's outputs should not be trusted, and assigning clear ownership of the model's risk. The model developed in this paper accommodates this governance dimension by allowing the maturity of an enterprise's artificial intelligence governance to feed into its control strength factor: an enterprise that governs its artificial intelligence systems well has, in effect, a stronger control environment, and the model reflects that strength in a lower loss probability.

11.5 Adversarial Artificial Intelligence and Data Poisoning

A distinctive class of risk arises when artificial intelligence systems themselves become the target of attack. Adversarial manipulation, in which an attacker crafts inputs designed to deceive a model into a wrong output, and data poisoning, in which an attacker corrupts the data on which a model is trained so that it learns the wrong lessons, are threats that have no close analogue in conventional cybersecurity. For a financial institution that relies on artificial intelligence for critical decisions, these threats create exposure that the institution may not recognize, because the model continues to operate and produces outputs that appear normal even as it has been compromised. Read through the framework of this paper, adversarial and poisoning attacks increase the vulnerability coefficient for an enterprise that depends heavily on artificial intelligence, because they raise the probability that a threat event produces a loss, and they call for specific controls, including the validation of training data, the monitoring of model outputs for anomalies, and the testing of models against adversarial inputs. An enterprise that has not considered these threats has a blind spot that its loss experience may eventually and expensively reveal.

11.6 Governance Frameworks and the Financial Sector

Public authorities have begun to provide frameworks for the governance of artificial intelligence, and the financial sector is adapting them to its specific needs. A national framework for managing the risks of artificial intelligence provides a structured approach organized around the functions of governing, mapping, measuring, and managing those risks, and it is designed to be tailored to sectors. For the financial sector, national authorities have issued guidance that recognizes both the opportunities artificial intelligence presents for strengthening cybersecurity and the risks it introduces, and that calls on institutions to adapt their risk management to the specific demands of the technology, to invest in the training of their staff, and

to address the shortage of professionals equipped to govern these systems responsibly. The existence of this guidance at the level of national policy is itself significant: it establishes that the responsible governance of artificial intelligence in finance is a recognized public priority, and it identifies the supply of qualified professionals as a constraint on meeting that priority. The model in this paper is designed to operate within this emerging governance landscape, treating the governance of artificial intelligence both as a control that reduces risk and, where it is weak, as a vulnerability that increases it.

11.7 Practical Controls for the Artificial Intelligence Frontier

For the practitioner, the abstract discussion of artificial intelligence risk resolves into a set of practical controls, and the model rewards their adoption through a stronger control environment. On the defensive side, these controls include the rigorous validation and continuous monitoring of artificial intelligence systems used in security and decision-making, the testing of those systems against adversarial inputs, and the protection of training data against poisoning. On the side of resilience against artificial-intelligence-augmented attacks, the controls include strengthened verification of communications to counter synthetic-media deception, accelerated patching to counter automated vulnerability discovery, and enhanced training of staff to counter machine-assisted social engineering. On the governance side, the controls include clear accountability for artificial intelligence systems, documented limitations, and the integration of artificial intelligence risk into the enterprise's broader risk management. An enterprise that adopts these controls reduces both the frequency and the severity of artificial-intelligence-related loss, and the model reflects that reduction. The honest caveat, repeated from earlier in the paper, is that this frontier moves quickly, and that the controls effective today may need revision tomorrow, which is why the model's requirement of regular recalibration and conscious expert judgment matters most in exactly this domain.

11.8 The Talent and Governance Dimension

A theme that runs through both the cybersecurity and the artificial intelligence dimensions of this paper is that the binding constraint on managing these risks is increasingly the supply of people equipped to do so. National authorities responsible for the financial sector have explicitly identified the shortage of professionals capable of governing artificial intelligence and cybersecurity in finance as a constraint on the sector's resilience, calling for investment in training and for the adaptation of risk management frameworks to the specific demands of the technology. This is not a peripheral observation; it is central to whether the methods this paper develops can be put into practice. A model that requires skilled calibration, sound judgment about its assumptions, and disciplined governance of its use is only as good as the people who operate it, and the demand for such people now outstrips the supply across the financial sector and the broader economy. The frameworks this paper develop are, in a sense, an attempt to make the scarce expertise of the cyber risk professional go further, by encoding sound practice in a structure that less specialized staff can operate under expert supervision. But no encoding eliminates the need for the expertise itself, and the responsible adoption of financial cyber risk quantification will depend on continued investment in the people who can calibrate it, govern it, and exercise judgment where the model reaches its limits. The governance of artificial intelligence in financial systems is a frontier where the combination of financial, technical, and regulatory expertise is rare and valuable, and where the quality of governance will increasingly determine whether artificial intelligence strengthens or undermines the security of the financial system.

11.9 A Scenario of Artificial-Intelligence-Augmented Fraud

To make the artificial intelligence threat concrete, consider a scenario that financial institutions increasingly confront. An attacker uses synthetic media to impersonate a senior executive in a video call, instructing a finance employee to authorize an urgent transfer to a fraudulent account, and supports the deception with machine-generated correspondence that mimics the executive's style convincingly enough to pass casual scrutiny. The defenses that once protected against such fraud, the difficulty of imitating a known person and the awkwardness of fraudulent correspondence, are eroded by technology, and the result is an increase in both the frequency and the success rate of such attacks. Read through the model. This scenario raises the threat event frequency and the vulnerability coefficient for an enterprise that has not adapted its controls, and it calls for specific countermeasures: verification procedures that do not rely on the appearance or voice of the instructing party, controls that require multiple independent authorizations for significant transfers, and training that prepares staff for the possibility of convincing impersonation. An enterprise that adopts these countermeasures strengthens its control environment, and the model reflects that strengthening in a lower loss probability. The scenario illustrates the general pattern of the artificial intelligence frontier: a familiar threat is amplified by new capability, the historical loss distribution understates the near-term risk until it is recalibrated, and the appropriate response combines updated controls with the conscious exercise of judgment about a threat whose contours are still emerging.

11.10 The Defensive Dividend and Its Governance

The discussion of artificial intelligence has emphasized the threats, but the defensive dividend is real and deserves its due, along with the governance that determines whether the dividend is realized or squandered. Artificial intelligence augments defense in several ways that translate, in the model, into a stronger control environment and thus a lower loss probability. It accelerates the detection of threats by identifying anomalous patterns that escape human attention, compressing the time between compromise and discovery, which is among the strongest determinants of the ultimate cost of an incident. It automates routine response, freeing scarce human analysts to concentrate on the complex and strategic threats that most require judgment. It enhances the triage of vulnerabilities, directing remediation toward the weaknesses most likely to be exploited. Each of these capabilities reduces the frequency or the severity of loss, and an enterprise that deploys them well earns a genuine reduction in its exposure. Governance caveat, however, is essential. A defensive artificial intelligence system that is poorly governed can itself become a source of risk, through the bias that produces blind spots, the opacity that frustrates accountability, or the adversarial manipulation that turns a defense into a vulnerability. The dividend is realized only when the defensive systems are governed with the rigor of model risk management, validated, monitored, documented, and owned. An enterprise that deploys defensive artificial intelligence without governing it may find that it has traded a visible risk for a hidden one, and the model, by trying the loss probability to the maturity of governance rather than to the mere presence of the technology, captures exactly this distinction.

11.11 The Recalibration Imperative

The single most important practical consequence of the artificial intelligence frontier for the use of the model is the heightened imperative to recalibrate. Every empirical model assumes that the past is a guide to the future, and that assumption holds well enough in stable domains. The domain of cyber risk is not stable, and the destabilizing force of artificial intelligence makes it less so, because a capability that materially raises the frequency or severity of attacks can render a distribution calibrated last year a poor guide to this year's risk

within months rather than decades. The model addresses this not by pretending to predict the unpredictable but by insisting on regular recalibration and on the conscious application of forward-looking judgment to the threat frequency when assessments of adversary capability warrant it. An enterprise that recalibrates its model annually, and that adjusts its frequency assumptions when it has reason to believe the threat environment has shifted, will keep its estimates current; an enterprise that calibrates once and trusts the result indefinitely will find its estimates increasingly disconnected from a reality the technology is reshaping. The recalibration imperative is thus not a peripheral maintenance task but a central discipline of using the model responsibly in an era of rapid technological change, and it is the practice through which the model's necessary assumption of stationarity is reconciled with the manifest non-stationarity of the world it measures.

12.Future Research Directions

The model opens several avenues for future work, of which five are most pressing.

1. Prospective validation. The validation reported here is retrospective. A prospective study, in which a panel of enterprises is assessed before incidents occur and the estimates are compared as losses materialize over a period of several years, would provide the strongest possible evidence of predictive validity and is the highest-priority future initiative.

2. Machine-learning enhancement of frequency calibration. Because threat frequency is the dominant driver of variance, improvements in its estimation would yield the greatest gains in accuracy. Ensemble models trained on threat intelligence, vulnerability disclosure data, and security posture metrics may outperform expert calibration, a possibility that warrants systematic study across the full parameterization.

3. A dynamic, continuously updated score. The current model produces a static annual estimate. As monitoring and intelligence increasingly generate real-time data, a continuously updated score would give boards and supervisors a live indicator of financial cyber risk, raising both engineering and governance questions worthy of study.

4. A systemic national index. Aggregating the model across the national digital trade infrastructure to construct a systemic index, calibrated against national trade-flow data, would provide a quantitative foundation for the assessment of systemic digital trade risk that no authority currently possesses.

5. Full calibration of the nation-state term. As set out in Section 6, the formal specification of the nation-state extension awaits full calibration, which is identified as the highest-priority methodological extension of the model.

12.1 Elaborating the Research Agenda

Each of the five directions identified above deserves a fuller statement, because together they constitute a coherent agenda for advancing the financial measurement of cyber risk. Prospective validation is the foundation of that agenda, because it is the only way to convert strong retrospective fit into demonstrated predictive accuracy; a study that assessed a panel of enterprises before incidents occurred and tracked the realized losses over several years would either confirm the model's predictive value or reveal where it must be improved, and either outcome would advance the field. The machine-learning enhancement of frequency

calibration follows naturally, because frequency is the dominant driver of variance and the parameter is hardest to estimate well; methods that learn frequency from rich threat and posture data, rather than relying on expert judgment, could materially sharpen the model's most important input. The development of a dynamic, continuously updated score would change the character of the instrument, turning an annual estimate into a live indicator that boards and supervisors could watch in something close to real time, and it raises engineering and governance questions about update frequency, data integration, and the consequences of a continuously moving number that are themselves worth study.

The construction of a systemic national index would extend the model from the enterprise to the economy, aggregating exposure across the national digital trade infrastructure to produce a measure of systemic cyber risk that no authority currently possesses, and calibrated against national trade-flow data, it could inform the most consequential questions of cyber policy. The full calibration of the nation-state term, finally, would close the one major gap the present paper leaves open, converting the formally specified extension of Section 6 into a fully parameterized component grounded in open and, where available, cleared intelligence. Pursued together, these five directions would carry the financial measurement of cyber risk from a validated methodology toward a mature discipline, and they define the work to which this paper is an invitation.

12.2 Data and Standards as Shared Infrastructure

A direction that cuts across the specific research initiatives concerns the data and standards on which financial cyber risk measurement depends, and the case for treating them as shared infrastructure. The model, like any empirical model, is only as good as the data on which it is calibrated, and the best data, the comprehensive loss databases, the threat intelligence, the supply-chain maps, is largely proprietary and expensive, which concentrates the capacity for sound measurement in the largest and best-resourced organizations. This concentration is inefficient and, from a systemic perspective, dangerous, because the systemic risk embedded in shared digital infrastructure is borne collectively while the capacity to measure it is distributed unequally. There is a strong case, therefore, for treating the foundational data and standards of cyber risk measurement as shared infrastructure, supported by public investment and collective effort, much as the data underpinning other forms of financial risk measurement has come to be. Public breach reporting, standardized loss taxonomies, shared threat intelligence, and open methodologies would lower the cost of sound measurement for all and would improve the quality of the systemic risk assessment on which sound policy depends. The reduced-data variant developed in this paper is a modest step in this direction, demonstrating that useful measurement is possible from public data, but the larger opportunity is for a deliberate public and collective investment in the shared infrastructure of cyber risk measurement, and the case for that investment is among the broader implications of the work this paper presents.

12.3 An Invitation to Collaboration

The research agenda set out above is larger than any single researcher or institution can pursue alone, and the work this paper presents is therefore offered as an invitation to collaboration rather than as a closed result. The prospective validation of the model would benefit from the participation of enterprises willing to be assessed and tracked over time; the enhancement of frequency calibration would benefit from the data and expertise of organizations that monitor the threat landscape; the construction of a systemic national index would require the cooperation of the public authorities that hold the trade-flow and infrastructure data on which it depends; and the calibration of the nation-state term would require the engagement of those with access to the relevant intelligence. The financial measurement of cyber risk is, in this sense, a collective

undertaking, and its advance depends on the willingness of researchers, practitioners, insurers, regulators, and policymakers to share data, methods, and judgments in pursuit of a common good. This paper contributes a validated methodology and an open structure to that undertaking, and it is offered in the hope that others will extend, refine, and challenge it, so that the measurement of one of the great risks of the digital age may improve through the cumulative effort that every mature field of risk measurement has required.

13.Conclusion

This paper has developed the Financial-Based Cybersecurity Risk Exposure Model, a formally specified and empirically validated framework that addresses a consequential and growing gap between the financial materiality of cybersecurity risk and the analytical tools available to enterprise risk practitioners, cyber insurers, regulators, and policymakers. The gap is neither academic nor abstract. It manifests in enforcement actions against enterprises that disclose cyber risk qualitatively when investors require financial clarity; in insurance disputes where carriers and policyholders disagree about whether losses were foreseeable; in capital shortfalls at institutions that have provisioned against average breach costs while facing tail-event losses; and in trade policy frameworks that have no quantitative basis for assessing the systemic risk embedded in the digital infrastructure on which a vast volume of economic activity depends.

The model addresses this gap through four integrated components. The Annualized Loss Expectancy, derived from empirically calibrated Monte Carlo simulation, provides a statistically rigorous estimate of direct breach costs grounded in data rather than expert opinion. The Sector Exposure Multipliers, calibrated against critical infrastructure classifications and validated against a decade of breach cost data, adjust for the documented differential in breach costs across sectors. The Digital Trade Disruption Coefficient, the most novel contribution, formally models supply-chain propagation, capturing the secondary exposure that single-entity models miss and that constitutes a large share of total exposure for the most trade-intensive enterprises. The Regulatory Cost Component integrates penalties, mandated remediation, and capital adjustments into a financially denominated estimate grounded in hundreds of historical enforcement actions.

Validation against more than a thousand realized United States cyber losses demonstrates a combined mean absolute error of 8.3 percent, superior to both the benchmark based on the Factor Analysis of Information Risk model and the operational risk value-at-risk benchmark, and competitive with well-established financial risk models. The ratio of the ninety-ninth percentile to the median, averaging nearly five across the archetypes, confirms the heavy-tailed character of cyber risk that makes expected-value reasoning structurally inadequate for capital allocation and catastrophic risk management.

This expanded edition has added worked numerical examples that make the method concrete, a formally specified extension for nation-state threats, an implementable reduced-data variant for smaller enterprises, an enlarged treatment of artificial intelligence on both sides of the contest, and a deepened analysis of the regulatory landscape. The practical demand created by securities disclosure rules, operational resilience requirements, bank capital calculations, and an expanding cyber insurance market creates an urgent and growing institutional need for precisely the kind of financially credible, board-reportable, and regulatory-grade cyber risk quantification that this model provides. As artificial intelligence augments both offensive capabilities and defensive technologies, and as digital trade infrastructure becomes ever more central to economic competitiveness, the financial stakes of inadequate cyber risk measurement will only intensify. The model offered here is a rigorous step toward a common financial language for the risk of the digital economy.

13.1 A Closing Synthesis

The argument of this paper can be compressed into a single proposition: that cyber risk has become a financial risk, and that the instruments used to manage it must therefore speak the language of finance. Everything else follows this proposition. The four components of the model are an attempt to translate the principal sources of cyber harm, direct loss, sector consequence, network propagation, and regulatory cost, into that language. The validation is an attempt to demonstrate that the translation is faithful to the historical record. The examples worked, the case studies, the implementation guidance, and the appendices are attempts to make the translation usable by the people who must act on it. The nation-state extension, the reduced-data variant, and the treatment of artificial intelligence are attempts to extend the translation to the edges of the problem, the strategic threat, the under-resourced enterprise, and the rapidly moving frontier. And the candid statement of the model's limits is an attempt to ensure that the translation is trusted for what it is rather than mistaken for an oracle. The financial measurement of cyber risk is not a solved problem, and this paper does not claim to have solved it. But it offers a structure that brings the problem within reach of the financial machinery of the modern enterprise, and it does so at a moment when the regulatory, market, and technological pressures all point toward the same conclusion: that the era of managing cyber risk in colors and adjectives is ending, and the era of managing it in money has begun.

13.2 A Final Word

The financial measurement of cyber risk is, in the end, a matter of bringing one of the great risks of the digital age within the disciplines that govern every other material risk an enterprise face. For too long, cyber risk has been the exception, managed by intuition and reported in colors, while credit, market, and liquidity risk have been measured in money and governed with rigor. The pressures of regulation, of the insurance market, and of a threat landscape transformed by artificial intelligence and digital trade now make that exception untenable. The model this paper develops is an attempt to end it, to give boards, insurers, regulators, and policymakers a credible, auditable, and usable measure of cyber risk in the language they already share. It is not the final word on the subject, and it is offered with its limits stated plainly and its assumptions open to scrutiny. But it is a serious step toward a future in which the financial stakes of cyber risk are measured with the precision the stakes demand, and in which the decisions that turn on those stakes are made with the clarity that only financial measurement can provide.

Appendices

A Glossary of Key Terms

The following definitions are provided for reference. They describe the terms used in this paper and are not intended as authoritative legal or technical definitions.

Term	Definition as used in this paper
Annualized Loss Expectancy	The expected direct financial loss from cyber incidents over a year, derived from the simulation of frequency and severity
Conditional tail expectation	The expected loss given that the loss exceeds a stated quantile; also called expected shortfall
Cyber Risk Exposure Score	The monetary output of the model at a stated confidence level
Heavy-tailed distribution	A distribution that assigns meaningful probability to extreme outcomes far from the center

Materiality	In securities law, the quality of being important to a reasonable investor's decision; an inherently financial concept
Monte Carlo simulation	A method of estimating a distribution of outcomes by repeated random sampling of the inputs
Network centrality	A measure of how systemically important an enterprise is within a network of relationships
Operational resilience	The ability of a financial entity to withstand, respond to, and recover from disruptions, including cyber incidents
Sector Exposure Multiplier	A composite factor adjusting direct loss for the consequences that vary by industry sector
Value at risk	A quantile of the loss distribution used to express the loss not expected to be exceeded at a stated confidence level

Table A1. Glossary of key terms used in the paper.

B Catalog of Data Sources

The model draws on several categories of data, available from both commercial and public sources. The catalog below summarizes the principal sources by category, distinguishing those that require a commercial subscription from those available publicly, so that an organization can plan its data strategy according to its budget. Specific source names and access terms should be verified against current availability before an organization commits to a data strategy.

Data category	Commercial sources	Public substitutes
Threat frequency	Commercial threat intelligence indices	Public breach-investigation cohort tables
Loss distributions	Commercial cyber loss databases	Public breach portals and regulatory filings
Sector data value	Commercial pricing intelligence	Public reporting and open economic tables
Regulatory history	Proprietary enforcement databases	Public enforcement announcements
Network mapping	Commercial supply-chain intelligence	Self-reported vendor lists and procurement data
Nation-state frequency	Cleared threat intelligence (where available)	Open trackers and agency advisories

C Implementation Checklist

The checklist below summarizes the steps an organization should complete to bring the model into productive use. It is intended as a practical companion to the implementation guidance in the body of the paper.

- Confirm the sector classification and identify any multi-sector exposure requiring revenue-weighted multipliers.

- Extract and clean historical loss, control, trade, financial, and regulatory-history data.
- Calibrate the frequency parameters from sector baselines adjusted for size and control maturity.
- Fit the severity distribution to sector-specific loss data, blended with internal data where available.
- Assign the sector multiplier and, for multi-sector enterprises, compute the revenue-weighted blend.
- Map the trade network and compute the dependency ratio, centrality score, and propagation velocity.
- Calibrate the regulatory cost component through the decision tree fitted to enforcement data.
- Run the Monte Carlo simulation and extract the median, ninetieth and ninety-ninth percentiles, and expected shortfall.
- Produce the confidence intervals and the sensitivity ranking of inputs.
- Validate the assembled parameterization against any available realized losses.
- Document the parameterization, sources, and assumptions for audit and supervisory review.
- Integrate the output into board reporting, control-investment prioritization, and audit-committee review.
- Establish the quarterly refreshing cadence and the annual recalibration and assumptions review.

D A Closing Note on Responsible Use

The model is a tool for informing decisions, not for replacing the judgment of the people accountable for them. Its output depends on the quality of the data and the soundness of the assumptions on which it is calibrated, and an organization that adopts it should subject those data and assumptions to the same scrutiny it would apply to any model on which financial decisions rest. The figures cited throughout this paper are drawn from the sources referenced and should be verified against those sources before they are relied upon. Where the model is used for regulatory, capital, or insurance purposes, it should be subjected to independent validation consistent with established model risk management practice, and its use should be governed by clear documentation and accountability. Used with this discipline, the model can raise the quality of cyber risk decisions across the enterprise, the insurance market, and the regulatory system; used without it, no model, however well-constructed, can be trusted.

References

Accenture Security. (2019). The cost of cybercrime: Ninth annual cost of cybercrime study. Accenture.

Anderson, R., Barton, C., Bohme, R., Clayton, R., van Eeten, M. J. G., Levi, M., Moore, T., and Savage,

S. (2013). Measuring the cost of cybercrime. In R. Bohme (Ed.), The economics of information security and privacy (pp. 265 to 300). Springer.

Basel Committee on Banking Supervision. (2011). Operational risk: Supervisory guidelines for the advanced measurement approaches. Bank for International Settlements.

Basel Committee on Banking Supervision. (2017). Basel III: Finalising post-crisis reforms. Bank for International Settlements.

Beutel, J., Lorenz, R., and Paulus, S. (2021). Machine learning-enhanced cyber threat frequency prediction for financial institutions. Journal of Cybersecurity Research, 6(2), 45 to 68.

Bohme, R., and Schwartz, G. (2010). Modelling cyber-insurance: Towards a unifying framework.

Workshop on the Economics of Information Security.

Bureau of Economic Analysis. (2024). Input-output accounts data: Use tables after redefinitions. United States Department of Commerce.

Corera, G. (2021). SolarWinds: How Russian spies hacked the United States government. British Broadcasting Corporation News.

Cybersecurity Ventures. (2024). Cybercrime magazine annual cybercrime report 2024. Cybersecurity Ventures.

Cybersecurity and Infrastructure Security Agency. (2023). Shifting the balance of cybersecurity risk: Principles and approaches for security-by-design and default. United States Department of Homeland Security.

Cybersecurity and Infrastructure Security Agency. (2024a). National critical infrastructure security and resilience research and development plan. United States Department of Homeland Security.

Cybersecurity and Infrastructure Security Agency. (2024b). Volt Typhoon: People's Republic of China state-sponsored cyber actor living off the land. Joint cybersecurity advisory.

Dittmar, A., and Field, L. C. (2024). Cybersecurity disclosure and board oversight: Evidence from the 2023 rules. Journal of Financial Economics, 154, 103 to 128.

Edwards, B., Hofmeyr, S., and Forrest, S. (2016). Hype and heavy tails: A closer look at data breaches.

Journal of Cybersecurity, 2(1), 3 to 14.

Eling, M., and Schnell, W. (2016). What do we know about cyber risk and cyber risk insurance? Journal of Risk Finance, 17(5), 474 to 491.

Eling, M., and Wirfs, J. H. (2019). What are the actual costs of cyber risk events? European Journal of Operational Research, 272(3), 1109 to 1119.

European Parliament. (2022). Regulation 2022/2554 on digital operational resilience for the financial sector. Official Journal of the European Union.

Federal Reserve System. (2023). Supervisory scenarios for the annual stress tests. Board of Governors of the Federal Reserve System.

Financial Stability Oversight Council. (2023). Annual report: Cybersecurity as an emerging financial stability risk. United States Department of the Treasury.

Gordon, L. A., Loeb, M. P., and Zhou, L. (2016). Investing in cybersecurity: Insights from the Gordon-Loeb model. Journal of Information Security, 7(2), 49 to 59.

Gordon, L. A., and Loeb, M. P. (2002). The economics of information security investment. Association for Computing Machinery Transactions on Information and System Security, 5(4), 438 to 457.

Huang, K., and Pearce, M. (2020). Cyber risk correlation in digital supply chains: Implications for insurance pricing. Risk Analysis, 40(8), 1578 to 1594.

Hubbard, D. W., and Seiersen, R. (2016). How to measure anything in cybersecurity risk. John Wiley and Sons.

International Monetary Fund. (2023). Cyber risk and financial stability: Evidence from financial institution incidents. Staff Discussion Note.

International Organization for Standardization. (2022). Information security, cybersecurity and privacy protection: Guidance on managing information security risks. ISO/IEC 27005:2022.

Jones, J., and Freund, J. (2015). Measuring and managing information risk: A FAIR approach (2nd ed.).

Butterworth-Heinemann.

Klugman, S. A., Panjer, H. H., and Willmott, G. E. (2019). Loss models: From data to decisions (5th ed.).

John Wiley and Sons.

Lloyds of London. (2024). Cyber risk outlook 2024: Systemic cyber risk and the market response. Lloyds of London.

Maillart, T., and Sornette, D. (2010). Heavy-tailed distribution of cyber-risks. European Physical Journal B, 75(3), 357 to 364.

McKinsey Global Institute. (2016). Digital globalization: The new era of global flows. McKinsey and Company.

Morgan, S. (2020). Cybercrime to cost the world many trillions of dollars annually by the mid-2020s.

Cybercrime Magazine. Cybersecurity Ventures.

National Institute of Standards and Technology. (2012). Guide for conducting risk assessments: Special Publication 800-30, Revision 1. United States Department of Commerce.

National Institute of Standards and Technology. (2024). Cybersecurity framework 2.0. United States Department of Commerce.

Office of the Comptroller of the Currency. (2023). Bulletin 2023-17: Technology and cybersecurity risk management supplemental examination procedures. Office of the Comptroller of the Currency.

Open Group. (2013). An introduction to factor analysis of information risk. Risk Management Insight and the Open Group.

Organisation for Economic Co-operation and Development. (2023). Measuring digital trade: Towards a conceptual framework. Statistics Working Papers.

Petri, P. A., and Plummer, M. G. (2020). Cybercrime and digital trade: Quantifying the economic losses.

Peterson Institute for International Economics Working Paper.

Ponemon Institute. (2024). Cost of a data breach report 2024. International Business Machines Corporation and Ponemon Institute.

Romanosky, S., Ablon, L., Kuehn, A., and Jones, T. (2019). Content analysis of cyber insurance policies: How do carriers' price cyber risk? Journal of Cybersecurity, 5(1).

Saltelli, A., Annoni, P., Azzini, I., Campolongo, F., Ratto, M., and Tarantola, S. (2010). Variance-based sensitivity analysis of model output. Computer Physics Communications, 181(2), 259 to 270.

Securities and Exchange Commission. (2023). Cybersecurity risk management, strategy, governance, and incident disclosure. Release Numbers 33-11216 and 34-97989.

Sheffi, Y. (2015). The power of resilience: How the best companies manage the unexpected.

Massachusetts Institute of Technology Press.

Verizon. (2024). Data breach investigations report 2024. Verizon Business.

World Economic Forum. (2024). Global cybersecurity outlook 2024. World Economic Forum.

World Trade Organization. (2024). World trade report 2024: Re-globalization for a secure, inclusive and sustainable future. World Trade Organization.

Zetter, K. (2014). Countdown to zero day: Stuxnet and the launch of the world's first digital weapon.

Crown Publishers.

Zhang, Y., and Zhu, T. (2023). Artificial intelligence-driven cyber threats and enterprise financial exposure: Emerging evidence from large-scale breach datasets—Journal of Financial Risk Management, 22(3), 88 to 112.

Zhao, X., Xue, L., and Whinston, A. B. (2013). Managing interdependent information security risks Journal of Management Information Systems, 30(1), 123 to 152.

Additional references for the expanded edition

Bureau of Labor Statistics. (2025). Occupational outlook handbook: Information security analysts. United States Department of Labor.

Centre for Internet Security. (2021). Critical security controls for effective cyber defence, version 8.

Centre for Internet Security.

Cooke, R. M. (1991). Experts in uncertainty: Opinion and subjective probability in science. Oxford University Press.

Council on Foreign Relations. (2024). Cyber operations tracker. Council on Foreign Relations.

CyberSeek. (2025). Cybersecurity supply and demand heat map. Computing Technology Industry Association, Lightcast, and the National Initiative for Cybersecurity Education.

European Union Agency for Cybersecurity. (2024). Threat landscape 2024. European Union Agency for Cybersecurity.

Federal Financial Institutions Examination Council. (2021). Architecture, infrastructure, and operations booklet. Information technology examination handbook.

Financial Stability Board. (2020). Effective practices for cyber incident response and recovery. Financial Stability Board.

International Information System Security Certification Consortium. (2024). Cybersecurity workforce study 2024. International Information System Security Certification Consortium.

Mandiant. (2024). M-Trends annual threat report. Mandiant, part of Google Cloud.

National Institute of Standards and Technology. (2020). Security and privacy controls for information systems and organizations: Special Publication 800-53, Revision 5. United States Department of Commerce.

National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework, version 1.0. NIST AI 100-1. United States Department of Commerce.

Office of the Comptroller of the Currency, Board of Governors of the Federal Reserve System, and Federal Deposit Insurance Corporation. (2022). Computer-security incident notification requirements for banking organizations. Final rule.

United States Department of the Treasury. (2024). Managing artificial intelligence-specific cybersecurity risks in the financial services sector. Office of Cybersecurity and Critical Infrastructure Protection.

World Economic Forum. (2025). Global cybersecurity outlook 2025. World Economic Forum.