Cybersecurity Ratings to Support Domain-Specific Financial Decisions: Applications to National Infrastructures, Cyberinsurance, Gordon-Loeb Model, and NPV Analysis

William Yurcik † ‡ [1] [0009-0004-8453-3898] Gregory Pluta [2] [0009-0004-4760-1441]

João Eduardo Aparecido Luisi Vieira [3] [0009-0002-6614-5811]

Fábio Roberto de Miranda [3] [0009-0004-1101-0631]

[1] Centers for Medicare & Medicaid Services (CMS) Baltimore MD USA
william.yurcik@cms.hhs.gov
[2] University of Illinois at Urbana-Champaign, Champaign IL USA
gpluta@illinois.edu
[3] Insper Institute of Education and Research, São Paulo, Brazil
joaoealv@insper.edu.br
fabiomiranda@insper.edu.br

Abstract— Cybersecurity ratings based on algorithmically combining weighted metrics provides a quantitative measurement of an organizational security posture. These cybersecurity measurements can then be (1) baselined for measuring continuous improvement, (2) used in comparison to peers in the same domain-specific context, (3) possible cybersecurity financial investment interventions considered, and (4) return-on-investment financial results quantitatively measured.

In this paper we use cybersecurity rating to first show practical results in the use of cybersecurity ratings for two large national infrastructures – all the hospitals in the U.S. and all the hospitals in Brazil. To our knowledge this is the first example of the use of cybersecurity ratings for cybersecurity financial decision-making to be designed and demonstrated for large national infrastructures.

This real-world practical experience leveraging cybersecurity ratings leads to theoretical discussion illustrating how cybersecurity ratings are currently being incorporated in cyberinsurance calculations, the Gordon-Loeb cybersecurity investment model, and the Net Present Value multiyear portfolio decision-making process.

Keywords: cyber-risk quantification, cyber risk management, cyber-defense strategy, cybersecurity decision-making, cybersecurity investment, cyber ROI

† Corresponding Author

‡ Organizational Disclaimer: "The views presented herein do not represent the views of the Federal Government."

1 Introduction – Current State-of-Affairs and a Way Forward

The current state of cybersecurity management circa 2026 is reactive solving problems as they arise.[1] The Cybersecurity & Infrastructure Security Agency (CISA) mandated security management for U.S. Federal agencies consisting of enterprise dashboards showing system vulnerabilities that have been identified but unpatched and/or otherwise not yet remediated [1][2]. Log-based security management (e.g. Splunk) and SIEM-based security management (Security Information & Event Management e.g. RSA NetWitness) consist of enterprise dashboards of prioritized alarms [3][4]. Compliance-based security management (e.g. FISMA) use an audit control checklist in comparison with a security standard (e.g. NIST 800-53); however, audit controls are not weighted such that one documentation finding is the same as one unimplemented technical control finding leading to the characterization of "check-the-box" [5][6][7]. Lastly, outsourcing security management to an external entity (Managed Security Services Providers) only transfers responsibility to contractual agreements without solving the reactive stance to solving problems as they arise [3][6][8]. The current state of cybersecurity reactive problem-solving is unsustainable, especially given recent advances in AI which attackers will undoubtedly incorporate in near-future attacks.

In contrast, proactive cybersecurity is a forward-thinking strategy focused on preventing cyberattacks by identifying, assessing, and fixing vulnerabilities before they can be exploited, rather than reacting after a breach occurs. It involves continuous activities like threat hunting, penetration testing, security awareness training, vulnerability management, monitoring, and patching to build a strong security posture, stop threats early, and reduce overall risk.

However, proactive cybersecurity also faces challenges like a rapidly evolving threat landscape, a shortage of skilled professionals, integrating complex automation/technology, overcoming "the human factor" (error/awareness), balancing resources between proactive/reactive, securing the expanding attack surfaces (cloud, IoT, supply chain), and gaining leadership buy-in and budget support for continuous adaptation.

The experience we share in this paper is that cybersecurity risk management can be quantified algorithmically with empirically-based cybersecurity ratings which can be leveraged to enable optimized practical and theoretical financial decision-making.

We first describe work toward achieving domain-specific SOC implementations and sharing initial empirical results from these implementations. Section 2 describes the underlying techniques we leverage in order to provide cybersecurity situational awareness on the scale needed for national infrastructures. Sections 3 and 4 provide background about the two large national infrastructures targeted for implementation of these techniques. Section 5 shares empirical financial data analysis in the form of baselines,

[1] This observation is based on the combined leadership experience of all the authors, one of whom has led multiple cybersecurity fusion integration centers (aka SOCs) of 150 analysts for over two decades.

scatterplots, and return-on-investment decision-making support. We next describe how cybersecurity ratings are currently be used as inputs to cyberinsurance calculations as well as theoretical input usage for the Gordon-Loeb cybersecurity investment model and Net Present Value portfolio decision-making analysis. We end with a summary and conclusions in Section 7.

2 Data Reduction via Metrics and Ratings

We focus on designing a SOC where cybersecurity posture can be empirically baselined (on a large scale) and then strategic interventions to improve cybersecurity posture can be measured with quantitative results (on a large scale). To further unpack scalability at a large scale, even if able to produce quantitative security measurements, and given automation support, the volume of security metric information at some point will become too large for human decision-making to comprehend relationships, interactions, and emergent properties when making strategic cybersecurity decisions.

There are two general techniques that can be leveraged to help address scalability. First, numerical data reduction techniques can combine data measurements from multiple sources while retaining underlying information [3]. Second, humans have extraordinary visual processing capabilities, especially for pattern recognition changes, capabilities estimated to be about 10 Mbps with brain reaction times on the order of 150ms [9].

In order to achieve scalable SOC management, we converged on a two-stage approach consisting of (1) numerical data reduction techniques to reduce data volume and (2) data visualization techniques designed to present information to human decision-makers. After initial proof-of-concept experiments and in-house trial-and-error adjustments, we are implementing this two-stage approach for complex real-world environments.

2.1 Cybersecurity Metrics

One of the most frustrating and ultimately dangerous things about cybersecurity is that you can *almost* measure it.[2] There are many component parts of an overall cybersecurity posture that can and/or should be measured. Composing and assessing overall cybersecurity posture from component parts is allusive, currently an unsolved problem, and may never be completely solved in a formal mathematical proof [10].

Nonetheless, there still remains a vital organization and engineering need to accurately assess overall security posture beyond subjective qualitative opinion. Unfortunately, misinformation and snake oil are also starting to fill this space. The work we present here is an attempt at assessing overall cybersecurity posture quantitatively while

[2] This observation is a paraphrase of what Matt Blaze insightfully expressed about cryptography, a data protection technique within cybersecurity [11].

acknowledging it is an approximation but dedicated to incremental improvement as indicated by incorporating peer review from rigorous research forums. Insistence on perfection should not prevent implementation of "good enough" improvements over the status quo, especially when a vital need exists.

It is beyond the scope of this paper to comprehensively discuss security metrics, but we will attempt a short concise overview. NIST defines a security metric as a useful measurement tool that can be used to support human decision-making toward improving cybersecurity performance [12]. Despite this simple intuitive definition there is nothing even close to a consensus identifying security metrics best practices, rather security metrics are uniquely dictated by enterprise environments and the responsible people in charge of cybersecurity.

This lack of consensus on identifying security metrics is not from lack of trying! Rather quite the opposite, there have been 20+ years of focused forums on the topic of security metrics dating back to NIST in June 2000. We provide an incomplete list of major security metric forums and reference important contributions outside of these forums [13-39]:

- NIST Computer System Security and Privacy Advisory Board (CSSPAB) *"Approaches to Measuring Security"*, June 2000
- *Workshop on Security Metrics (MetriCon)* 2006-2019
- *International Workshop on Security Measurements and Metrics (MetriSec)* 2010-2012
- *International Workshop on Quantitative Aspects in Security (QASA*) 2012-2017

In Figure 1 we categorize examples of potential proactive cybersecurity metrics which can then be measured and used to quantitatively describe the cybersecurity posture of a system in as few measurements as possible.

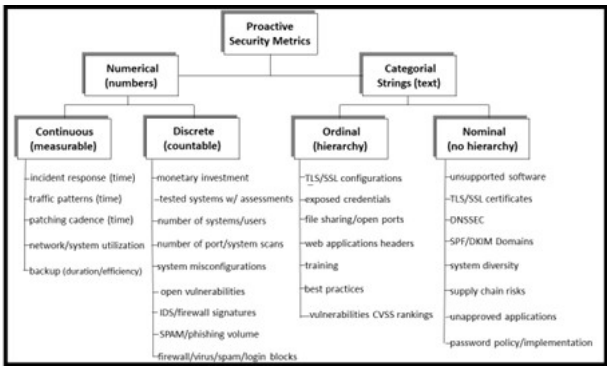

Figure 1. Categorization of Potential Cybersecurity Metrics

Note these proposed metrics look forward beyond reactive dashboard tracking the remediation of Common Vulnerabilities and Exposures (CVEs) and Known Exploited Vulnerabilities (KEVs) [40,41]. The objective for these security metrics is to provide indication what may happen next, beyond what has already happened. The security metrics in Figure 1 can all be measured/quantified in different ways from numerical-native metrics such as incident-response-times and number-of-tested-systems-with-assessments to categorical string-native metrics that can be quantified in rankings (different levels of reported exposed credentials) or binary (existence of unapproved applications).

Figure 2 shows the 13 metrics we have selected as inputs to a cybersecurity ratings algorithm. These 13 metrics are empirically measured using non-intrusive, external, and automated methods, including scanning public internet data, analyzing botnet traffic, tracking file-sharing activity, and reviewing public breach disclosures. uses headless browsers to scan for web application vulnerabilities, SSL/TLS configurations, insecure ports, botnet traffic, and malware-infected systems. These inputs are normalized and updated regularly before being incorporated into a ratings algorithm. Since this aspect is not the focus of this paper, for a more in-depth description of our metric data gathering see [42].

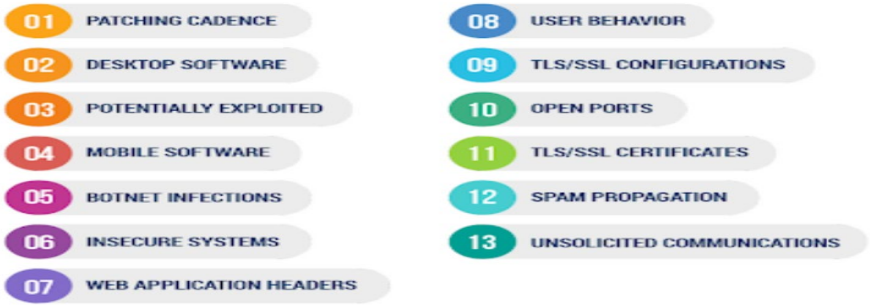

Figure 2. Selected Cybersecurity Metrics to be Empirically Measured

Each of these metric measurements is weighted for severity based on their correlation to security breaches, with higher-risk, more critical factors having higher weight contributions to the overall rating. The largest weight (70.5%) measures 11 different underlying sub-metrics for best practice implementation [patching cadence, web application headers, TLS/SSL certificates/configurations [42]. The next largest weight is an indication of compromised systems (27%) which measure evidence of preventing (or lacking to prevent) malicious or unwanted software [unsupported software, potentially exploited systems, botnet infection, insecure systems, spam]. The smallest weight is user behavior (2.5%) which measures three different activity metrics [open ports, password re-use, and file sharing traffic [42]. Since this aspect is not the focus of this paper, for a more in-depth description of our cybersecurity rating algorithm see [42].

For examples of the proactive nature of just two potential cybersecurity metrics, a shorter patching cadence has been documented to be correlated with less risk since it reduces the window of time that a system is vulnerable to a known exploit [43] and implementation of any or all of the following email-related protocols – the Sender Policy Framework (SPF) protocol, the DomainKeys Identified Mail (DKIM) protocol, and the Domain-based Message Authentication, Reporting & Conformance (DMARC) protocol - have proven effective at preventing email spoofing, reducing spam and potential for phishing attacks by verifying legitimacy of email senders [44].

Our solution is to combine our selected weighted security metrics in an algorithm to result in a single cybersecurity rating number.

2.2 Cybersecurity Ratings

Ratings are one of society's core tools for making decisions under uncertainty. We use them everywhere because they let us compress complexity into something comparable, scalable, and actionable. Examples include credit scores for creditworthiness; financial instrument bond ratings for pricing, eligibility, probability of default; vehicle crash test ratings for car safety; health risk factors for lifespan management; and stock market indexes to summarize financial market performance. These rating systems all have these attributes in common:

- Summarize complex information
- Enable comparison at scale
- Support threshold-based decisions
- Guide resource allocation
- Reduce uncertainty

Cybersecurity ratings - the result of combining weighted security metrics into an algorithm to result in a single number - are a powerful numerical data reduction technique which provides an indication of security effectiveness which has been validated against actual cybersecurity attacks [45]. Empirical research has found that externally observable cybersecurity ratings and related security posture metrics are statistically associated with the likelihood of experiencing data breaches, supporting their use as probabilistic risk indicators [46].

A low cybersecurity rating indicates increased likelihood of certain cyber incidents, not certainty of a breach, and should be used to prioritize investigation and remediation rather than to predict inevitable failure. Just like the predicted probability of rain in a weather forecast does not mean it will rain with certainty, a low rating does not mean that a breach is imminent, that defenses have failed internally, that sensitive data will be lost, that the organization is "unsafe", or that a specific attack will occur. Many organizations with low ratings never experience a breach, while some with high ratings do.

While ratings may provide a probability for organization risk to cyber incidents, cyber incidents are not deterministic, breaches require all the following chain elements and ratings only speak to one element of this chain:

1. Threat Actor
2. specific capability usable by a Threat Actor
3. Internet-exposed vulnerabilities *(cybersecurity ratings provide an empirical summary of this noting not all exposed weaknesses are exploitable in practice)*
4. Threat Actor detection of Internet-exposed vulnerability target(s) matching threat actor capability
5. Threat Actor selecting reachable Internet-exposed vulnerability target(s) to exploit and succeeding operationally

While cybersecurity ratings are useful for predicting directional risk reduction and relative improvement, they do not provide information about monetary loss and breach severity. For instance, cybersecurity ratings do not distinguish which systems support prioritized revenue streams, which processes are mission critical, which system outages are existential versus tolerable, data volume, and/or protected data from unprotected data. The best example of the limits to cybersecurity ratings is that two organizations with identical ratings may face wildly different consequences from the same incident. Ratings say *how likely* something is, not *how much it hurts*.

Most importantly, a given cybersecurity rating number is only one data point in time. A single rating shows where you are; ratings over time show how you behave—and behavior is what predicts risk. Because risk is dynamic, not static, a single rating is only a snapshot; ratings over time reveal behavior, trajectory, and resilience—which are far more predictive and useful for decisions - the longitudinal trend of a rating based on where it was in the past toward where it may likely be going in the future. A longitudinal rating over time reveals whether risk is improving vs deteriorating or volatile/fragile vs stable. Stable improvement is more trustworthy than a high but erratic rating. Sustained performance over time suggests repeatable processes, operational discipline, and control durability. One-time spikes often reflect temporary fixes, superficial remediation, and “score optimization” without structural change.

As an example, Figure 3 shows a longitudinal cybersecurity rating line over time with 13 significant annotated events, starting with an initial increasing/improving gradient followed by stable sideways trend, ending in a steep decline with an event indicating fragile volatility. Volatility itself is treated as risk.

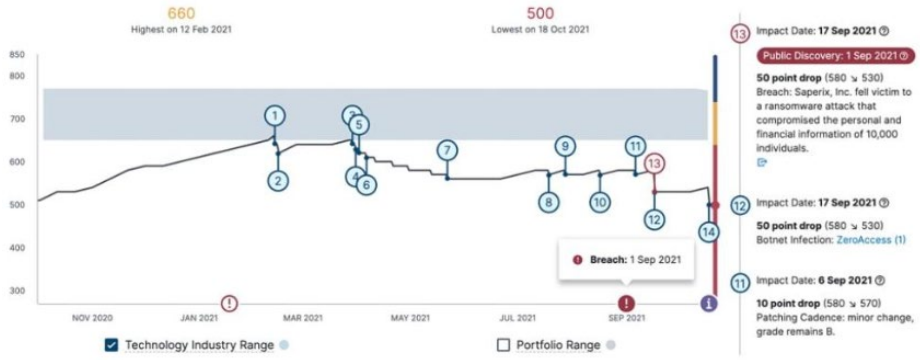

Figure 3. Annotated Longitudinal Cybersecurity Rating Line

2.3 Supplementing Cybersecurity Ratings with Context: Threats to Healthcare

While cybersecurity ratings may be analogous to a general health evaluation, then industry-specific threat planning is analogous to knowing whether the patient is an athlete, a cardiac patient, or a diabetic - the treatment depends on the context! While cybersecurity ratings communicate an organization's general exposure; industry-specific threat planning tells you *how*, *why*, and *where* an organization is most likely to be attacked—and what would matter most if attacked. Examples of how threat planning is important to consider for different domains include the following:

- Industry sectors attract different attacker motivations which ratings alone cannot capture.
- The same vulnerability has different consequences for different industry sectors.
- Attack techniques vary by industry sector - each industry faces distinct dominant attack patterns.
- Regulatory and legal exposure is industry-specific.
- Operational tolerance differs by industry - downtime tolerance varies dramatically.
- Industry context improves interpretation of cybersecurity ratings. The same rating may indicate something entirely different in another industry context.

While cybersecurity ratings provide a valuable, cross-industry signal of relative exposure and hygiene, industry-specific threat planning is essential because attacker motivations, attack methods, operational impact, and regulatory consequences vary dramatically by sector.

Cybersecurity threats specific to the healthcare domain include ransomware attacks, phishing, malware, and vulnerable/unpatched IoT medical devices, legacy systems, third-party vendors. Impacts include patient safety risks from disruptions to continuity

of care, exfiltration of legally-protected patient data for identity theft for fraudulent transactions and financial/reputational damage from public awareness of compromise. The potential impact of mortality in the healthcare domain makes it unique.

Cybersecurity threats in the healthcare domain have reached unprecedented attack levels, with the largest cybersecurity incident of any kind being the 2024 Change Healthcare nation-wide pharmacy disruption (impacting 190 million people) highlighting the vulnerability of healthcare infrastructure [47]. Healthcare remains a primary target because of its need for continuous operations and the value of medical data on the dark web.

Cybersecurity defense for healthcare can be improved by adopting domain-specific techniques such as specialized tools - like domain-specific web-application firewalls (WAFs) for healthcare web apps and/or domain-specific endpoint detection and response (EDR) for healthcare endpoints, applying least privilege within the healthcare domain, hardening domain-specific configurations, automating domain-specific updates, and continuous training tailored to healthcare domain threats (e.g., continuity of care, medical data, medical device IoT protocols). This approach builds layered defenses, reduces the overall attack surface, and aligns security with healthcare-specific business goals.

3 Domain Context: Healthcare in the U.S.

We are designing a SOC to monitor cybersecurity ratings for all of U.S. healthcare as a national infrastructure. Healthcare includes all organizations, people, and actions whose primary intent is to promote, restore, and/or maintain health. This includes all medical providers, out-patient urgent care, community clinics, nursing homes, specialized medical equipment providers, health insurers, the pharmaceutical industry, and many different types of hospitals.

USA healthcare covers a current population of 333 Million people, with private group insurance plans covering about 66% of the population, Medicaid covering 89 Million, Medicare covering 64.5 Million, the Affordable Care Act covering 21 Million, and 26 Million people with no health insurance.[3] In 2022, USA healthcare expenditure accounted for $4.5 trillion which is 17.3% of the U.S. GDP[4].

[3] As of May 2022 exactly 64,553,288 people were enrolled in Medicare and exactly 88,978,791 people were enrolled in Medicaid and Children's Health Insurance Program (CHIP). About 12M individuals are dually eligible for both Medicare and Medicaid, so are counted in the enrollment figures for both programs. In January 2024 the Affordable Care Act's Health Insurance Marketplace reached 21M for the 2024 plan year. In September 2023, the U.S. Census reported that for 2022 the number of uninsured U.S. citizens reached a record low of 26M or 7.9%. Note due to significant overlaps in coverage these numbers do not add to the current USA population for the year of study [48].

[4] In 2022, National Healthcare Expenditure (NHE) grew 4.1% to $4.5 trillion, or $13,493 per person, and accounted for 17.3% of Gross Domestic Product in 2022 [48].

U.S. healthcare is particularly challenging since the U.S. healthcare is a mixed economic system combining individual out-of-pocket payments, private health insurance (primarily linked with employment), and publicly-funded government health insurance (Medicaid and Medicare) where healthcare assets are both private and publicly-owned and prices are set by both supply-and-demand and regulatory fiat.

To tangibly assess the security posture of the mixed USA healthcare system we converged on hospitals as the central point touching every part of the industry – most providers have hospital privileges and hospitals are typically the parent organization of subsidiary activity such as associated out-patient services/facilities. We used multiple open-source authorities to assemble a database of 7,490 USA hospitals hosted at the University of Illinois at Urbana-Champaign which has been vetted multiple times.

A hospital is state-licensed institution whose function is to provide diagnostic and therapeutic patient services for medical conditions, with organized physician staff and registered nurses. The functional hospitals we are tracking include general hospitals, Short-Term Acute Care Hospitals (STACH), Long-Term Acute Care Hospitals (LTACH), Inpatient Rehabilitation Facilities (IRF), Skilled Nursing Facilities (SNF), short stay hospitals, behavioral hospitals, psychiatric care hospitals, children's hospitals, women's hospitals, teaching hospitals, and specialty care hospitals. Formal categories of hospitals include Acute Care/Critical Access Hospitals (ACH, fewer than 25 in-patient beds and greater than 35 miles from the next nearest hospital) and Safety-Net Hospitals (designated by the proportion of charity care provided). Figure 4 shows all USA hospitals mapped to their geographical coordinates in the continental USA.

Figure 4. USA Hospitals Mapped to Geographical Coordinates (hospitals in Alaska & Hawaii not included in this display)

We subdivided all U.S. hospitals into five separate categories in order to share more detailed analysis: (1) Indian Health Service (IHS) Hospitals, (2) Veterans Health

Administration (VHA) Hospitals, (3) Defense Health Agency (DHA) Hospitals, (4) Interstate Hospital Systems, and (5) Intrastate Hospital Systems.

4 Domain Context: Healthcare in Brazil

In combination, we are also designing a SOC to monitor cybersecurity ratings for all of U.S. healthcare as a national infrastructure. Brazil has the world's largest national healthcare system in several dimensions - as measured in land area coverage (3.3 million square miles), number of affiliated treatment centers (50,000 clinics and treatment centers) and user base (220 Million). Brazilian healthcare is characterized by a unique blend of public and private sector involvement as defined by two primary components: the Unified Health System (Sistema Unico de Saude aka SUS) and a substantial private healthcare market. Brazil is also diverse, with urban regions having healthcare on par with top world-class healthcare facilities, while remote areas struggle with access to basic care. Brazil's SUS provides universal healthcare access to over 220 million citizens [49-51]. The SUS system operates on a decentralized governance model centered on the Family Health Strategy (FHS) which coordinates care across the nation.

Private entities offer healthcare services to roughly 25% of the population, primarily through employer-sponsored insurance plans and out-of-pocket payments. There is also a significant presence of non-profit organizations, such as the Santas Casas, which operate hospitals that provide care funded by the public system. Entities that are private and for-profit can also perform services for SUS, creating a complex interplay between public funding and private delivery. The volume of healthcare providers makes securing all of the infrastructure challenging. Brazil has approximately 6,000 hospitals and 310,000 healthcare providers registered with the National Registry of Health Facilities [52], distributed across its 26 states and the Federal District.

Table I provides a breakdown of healthcare facilities by type and sector. Figure 5 shares the sectors and geographical distribution of hospitals across different regions of Brazil.

TABLE I
FACILITY TYPE VS SECTOR IN BRAZIL

Sector → / Type ↓	Private	Public	Total
Hospital	4108	1940	6048
Outpatient	277537	7202	284739
Others	28254	26259	54513
UBS	837	37087	37924
Total	310736	72488	383224

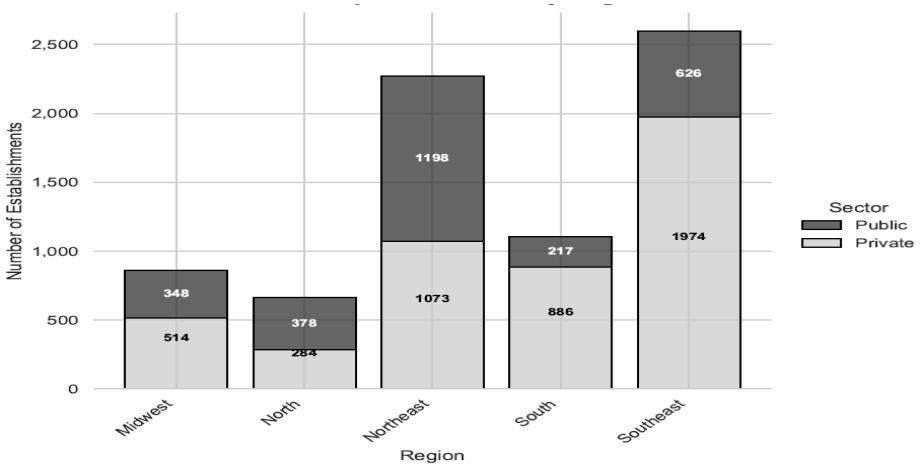

Figure 5. Distribution of Hospitals in Brazil by Geographical Region and Sector

Hospitals are a cornerstone of the Brazilian health system and are often a primary target of cybersecurity attacks. The Southeast region has over 2,500 hospitals, accounting for more than 40% of the total in the country [53]. This region includes populous states like Sao Paulo and Rio de Janeiro, which have a high concentration of both public and private healthcare facilities. The Northeast region has approximately 2,000 hospitals [53].

5 Cybersecurity Financial Management

Baselines provide a starting point for measuring continuous improvement as reflected in cybersecurity ratings. Achieving higher cybersecurity ratings will not happen on its own but will require strategic interventions/investments in order to maintain and improve [2,3,7]. On the flip side, without strategic cybersecurity interventions/investments over long periods of time decreased cybersecurity ratings will result as technology advances and existing security protection techniques degrade and become obsolete [2,3,7].

5.1 Cybersecurity Rating Baselines

Table II provides a comparison of the cybersecurity rating baselines for each of the five hospital systems identified which encompass 69% of all the hospitals in the U.S. The baselines of the IHS and VHA hospital systems are statistically significantly different from each other and statistically significantly different from both Interstate/Intrastate Hospital Systems since their 95% confidence intervals for their means do not overlap. However, the baselines of Interstate/Intrastate Hospital Systems are not statistically significantly different from each other since their 95% confidence intervals for their means overlap. This makes intuitive sense since both the IHS and VHA Hospital Systems have their own unique centralized IT coordination while Interstate/Intrastate Hospital

Systems each consist of many different independent hospital systems, with each hospital system acting independently with little, if any, IT coordination between hospital systems.

TABLE II
COMPARISON OF CYBERSECURITY RATING STATISTICS FOR FOUR DIFFERENT U.S. HOSPITAL SYSTEMS AND INDEPENDENT U.S. HOSPITALS

Security Rating Statistics	IHS	VHA	Interstate Hospital Systems	Intrastate Hospital Systems	Independent Hospitals
Mean	719.78	753.78	682.72	699.34	733
95% CI	+/- 7.25	+/- 2.96	+/- 12.00	+/- 5.62	na
Median	730	760	690	710	740
Range	650-760	690-780	500-800	460-800	550-800
Targets	12	25	50	29	na

5.2 Cybersecurity Rating Scatter Plots – Healthcare Domain-Specific

We have designed domain-specific scatterplots as a valuable SOC tool to identify possible hospital targets for further inspection, investigation, and possible intervention. Further inspection and investigation to determine the root cause of poor cybersecurity rating values and/or significant cybersecurity rating changes. Possible intervention is that if it can be determined that something can be done to improve cybersecurity protection, as reflected in improved cybersecurity ratings, then it will be considered.

Figure 6 shows the previous scatterplots of hospital systems we have analyzed with each scatterplot mapping cybersecurity ratings (vertical y-axis) versus hospital size as measured by inpatient hospital beds (horizontal x-axis). These scatterplots (projected on SOC tiled walls) represent approximately 2/3rds of all the hospitals in the U.S. - with the cybersecurity ratings for each hospital in a hospital system combined to be represented by one data point for each hospital system. The bottom leftmost Figure 6 scatterplot shows the distribution of cybersecurity ratings for hospitals in USA Interstate Systems (126 systems, 2,612 hospitals). The bottom rightmost Figure 6 scatterplot shows the distribution of cybersecurity ratings for hospitals in USA Intrastate Systems (523 systems, 2,571 hospitals).

The last row in Table II indicates the number of potential target hospitals/systems which would be the best candidates for further inspection and possible intervention, an intervention designed to result in a statistically significant increase in cybersecurity rating.

Graphically in Figure 6 we consider two dimensions for selecting hospitals for interventions/investment resulting in the largest beneficial patient outcome and the largest increase in security rating score (the largest hospitals with the lowest cybersecurity

rating) - basically hospital systems in proximity to the lower right quadrant of the scatter plot.

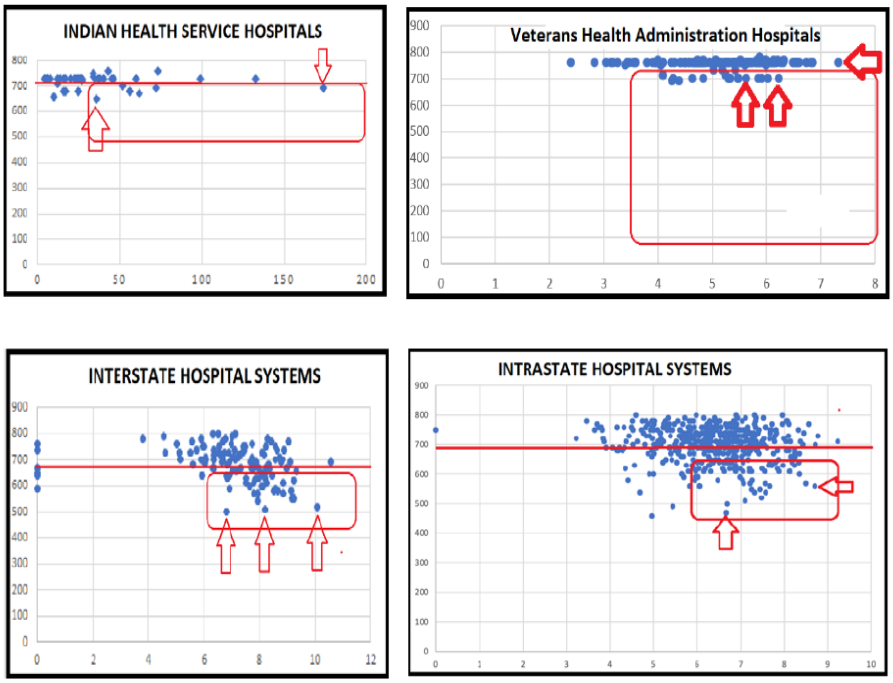

Figure 6. U.S. Hospital System Targets for Possible Intervention

In Figure 7 we show another scatterplot example, this time for independent rural hospitals (2,005 hospitals) looking to identify targets for further inspection and possible intervention. The horizontal x-axis is hospital size as measured by inpatient beds and the vertical y-axis are cybersecurity ratings. With red arrows we identify three independent hospitals with low cybersecurity ratings, two of the rural hospitals identified for possible intervention have about 200 inpatient beds and one of the hospitals identified for potential intervention is rather large with 600 inpatient beds.

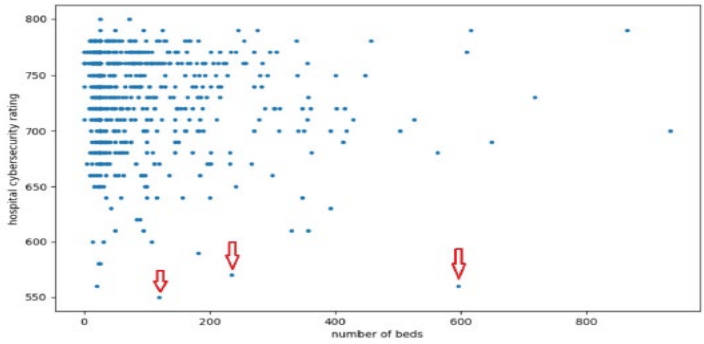

Figure 7. Independent Hospitals Targets for Possible Intervention

5.3 Measuring Domain-Specific Return-on-Investment (ROI) with Cybersecurity Ratings

An ROI decision is a financial choice to invest in a project based on its expected profitability, calculated by comparing the gain (net profit) to the cost, expressed as a percentage, to see if it generates sufficient value compared to other options. For our healthcare domain-specific SOC, we would like to make the most effective investments to increase cybersecurity ratings in hospital systems and independent hospitals for the most patients (as measured by inpatient beds). Calculating the cybersecurity rating movement versus cost of an intervention/investment will provide a Return-on-Investment (ROI) calculation that can be used for strategical planning [3,7].

$$ROI = (Net - Profit/Cost - of - Investment)x100\%$$

ROI Decision-Making:

1) Calculate the Potential ROI using the formula above to find the expected return.

2) Note that in our domain-specific case Net-Profit is the expected monetary gains expected from the investment in terms of cybersecurity attack damage avoided.

3) Cost of Investment: The initial outlay plus any associated expenses like maintenance and/or training.

4) Compare Options: Evaluate the calculated ROI for different possible investment projects.

5) Make the Choice: Select the investment with the highest positive ROI that meets your risk tolerance.

ROI is important because it shows the value of how much money is made (or payments avoided) for every dollar invested. ROI also allows direct comparisons between diverse investments that may otherwise be hard to compare. Ultimately ROI calculations help enterprises invest wisely in projects that offer the largest financial gains.

We consider two hypothetical security intervention/investment scenarios. Scenario One is a small ratings impact but broad intervention across a large number of hospitals based on three low-weighted security metrics in our cybersecurity ratings algorithm. This is a low intensity effort in resources at each hospital but treating more hospitals. Depending on the low-level treatment required at each hospital, it may be possible to accomplish treatment remotely via conferencing and shipment of equipment as needed.

Scenario Two is a large ratings impact but focused intervention involving a small number of hospitals who are performing poorly in security management. Prioritizing hospitals starting with the lowest security rating and working upward intervening to bring each treated hospital up to the highest cybersecurity rating prior to intervention. This is an intensive effort in resources at each hospital but treating less hospitals and less travel. Since this is a high level of treatment at each hospital, it cannot be accomplished remotely and will demand more time at each hospital.

Table III shows results from the two scenarios. The intervention of Scenario One, intervening by investing in a small treatment across a large number of hospitals results in only one hospital system (IHS) increasing its mean ratings with statistical significance (after interventions at 31 hospitals). The intervention of Scenario Two, intervening by investing in a large treatment across a small number of hospitals, results in all four hospital systems increasing their mean ratings with statistical significance. For these two scenarios, and an infinite number of other scenarios, the intervention/investment cost (or lack of intervention/investment) can be measured in security ratings changes and an ROI calculated. Intervention/investments can then be optimized for ROI, under a budget constraint, for evidence-driven strategic security management decisions.

TABLE III
SCENARIOS ONE/TWO STRATEGIC INTERVENTION RESULTS

Intervention Scenarios	IHS	VHA	Interstate Hospital Systems	Intrastate Hospital Systems
Scenario One "broad-and-shallow"	YES-31	NO-21	NO-41	NO-85
Scenario Two "focused-and-deep"	YES-7	YES-9	YES-12	YES-18

6 Cybersecurity Ratings as Inputs: Cyberinsurance, Gordon-Loeb Model (GL), and Net Present Value (NPV)

Organizations need to continually invest in maintaining, upgrading, and deploying innovative technologies to support their enterprise business mission given the dynamic nature of IT services. Failure to do so will expose an organization to increasing risks from software/system vulnerabilities, especially at the present time given new cyber threats being introduced with the use of AI. However, such investments have a budget constraint such that decisions must be made how best to spend limited money to reduce the most risk.

6.1 Cybersecurity Ratings as Input to Cyberinsurance

Cybersecurity ratings are currently widely used by cyber insurers in calculating premiums based on estimated relative likelihood of loss events segmented into risk tiers [59]. Cyberinsurance is a mechanism for transferring financial cyber risk from an organization to an insurer in exchange for predictable annual premiums, thus reducing volatility and protecting against unexpected financial losses an organization cannot absorb [54-58].

Cybersecurity ratings affect cyber insurance premiums but are not guarantees of coverage [59]. Before issuing a quote in pre-quote screening (risk triage), insurers decide two things: (1) is this risk insurable? and (2) how much underwriting effort is justified? Cybersecurity ratings help insurers quickly screen thousands of applicants, identify high-risk outliers, and decide whether to auto-quote, require in-depth underwriting, or decline. This is directly analogous to how credit scores are used before loan underwriting.

Cyber insurers do not rely on cybersecurity ratings alone [60]. They combine ratings with questionnaires, claims history, and domain-specific industry knowledge (revenue norms and threat landscape). Ratings contribute primarily to likelihood estimation, external attack surface risk, and empirical validation of self-reported cybersecurity assessments but not loss impact estimation.

Once underwriting is underway, cybersecurity ratings help insurers segment applicants into risk tiers and adjust premiums, retentions (deductibles), and coverage limits. For illustration:

- Higher rating → lower expected frequency → lower premium
- Lower rating → higher frequency → higher premium or exclusions

For applicants with low cybersecurity ratings, insurers may require remediation and an independent assessment (before binding), exclusion of specific attack types, and higher deductibles [62]. Many insurers now use continuous monitoring of longitudinal ratings after binding to monitor security posture drift, identify emerging risks, and in rare cases for mid-term adjustments.

Cyber insurers trust cybersecurity ratings because they correlate with certain claim types and they support risk differentiation at scale [60]. They allow insurers to maintain uniform consistency, external empirical validation, and reduce manual expert review costs. This directly mirrors how FICO credit scores are used in lending, telematic apps are used in auto insurance, and health risk scores are used in medical underwriting.

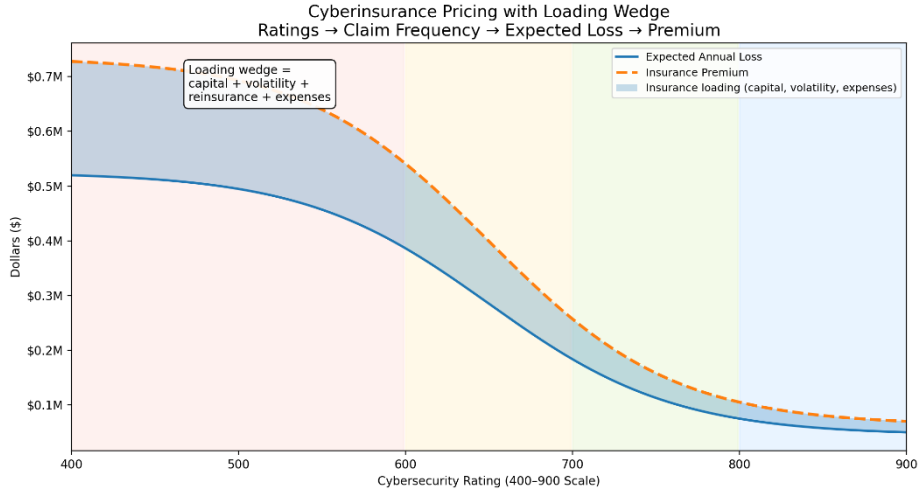

Figure 8. Cyberinsurance Premiums vs Expected Loss Curves

Figure 8 is a mathematical visualization showing how cybersecurity ratings are used in cyber-insurance pricing and underwriting. The axis indicates the cybersecurity rating (400–900 scale) and the Y-axis, denoted in Dollars ($), indicates the expected annual loss and the resulting premium. What immediately sticks out is that while the insurance premium tracks expected loss it sits higher. The insurance premium is higher than expected loss because insurers must charge more than the statistical average loss to cover uncertainty, operating costs, capital risk, and profit. That difference is called the insurance loading.

The expected loss curve (solid line in Figure 8) is the *actuarial average* of how much loss should be expected per year.

$$\textit{Expected Loss} = \textit{Claim Frequency} \times \textit{Claim Severity}$$

While the *expected loss* is an average across many domain-specific organizations, the premium must be higher than to account for *unexpected losses* such as attack clusters and correlated events so insurers charge extra (aka loading wedge) to absorb

unexpected losses. Cyber insurers have operating costs, reinsurance costs, and must also maintain capital to satisfy regulators and maintain credit ratings. Combined these operating costs, reinsurance costs, and capital requirements are all reflected in premiums.

$$Premium = Expected\ Loss \times (1 + Loading)$$

While cybersecurity ratings affect claim frequency and risk tiering, they do not determine operating expenses, capital requirements, and reinsurance costs as reflected in the loading wedge. The insurance premium curve sits above expected loss because insurers must price for uncertainty, capital risk, operating costs, and profit—not just average losses. While improving a cybersecurity rating lowers expected loss frequency, it does not eliminate the loading wedge [61].

Cybersecurity ratings correlate with breach event frequency, not with severity as measured in the cost for breach losses [45]. Insurers price in risk tiers based on severity, using underwriting and eligibility thresholds. Cyberinsurance premiums are nonlinear in tiers because insurers are pricing tail risk, correlation, and capital constraints, not just average loss—and pricing tail risk, correlation and capital constraints are not linear as security improves. Insurers are pricing to survive the worst cases.

Cybersecurity ratings do not include ***systemic risk*** [59, 63]. Thus, while improving cybersecurity ratings reduces average claim frequency for an organization, they do not reduce worst-case systemic losses which could lead to insurer insolvency. Cyber losses follow a heavy-tailed distribution in which many insured have no claims, but a few insured have events with extremely large losses which dominant severity calculations For example, one vulnerability can affect thousands of insureds and/or one cyber-attack campaign can trigger many simultaneous claims.

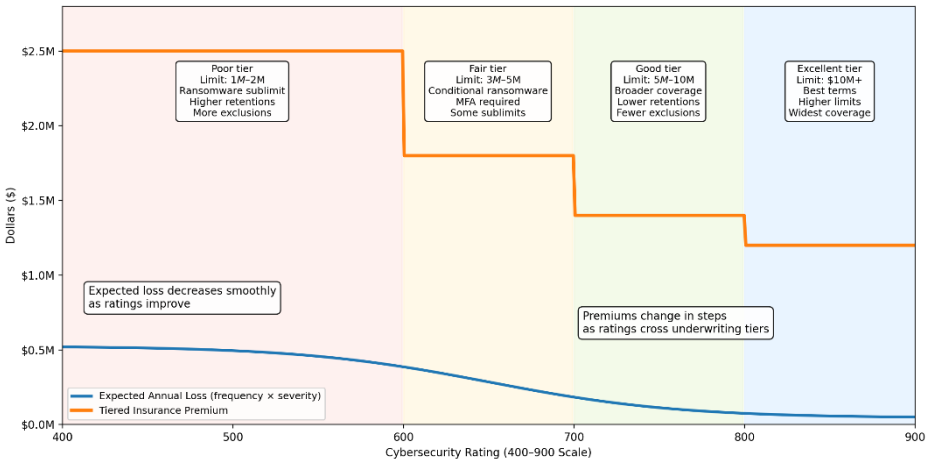

Figure 9. Cybersecurity Ratings Improvement vs Cyberinsurance Premium Changes

Figure 9 shows that while cybersecurity rating improvements reduce expected loss frequency smoothly as reflected in the expected annual loss line (lower line), insurance premiums fall in a step function (upper line) because insurers price by risk tiers, not by marginal improvements.

Thus, cybersecurity ratings affect insurance primarily through coverage availability and terms, while cyberinsurance premium levels adjust only in a step function at risk tiers [59]. As cybersecurity ratings improve, insurers will expand coverage limits and remove exclusions in discrete tiers (which is often more valuable than premium reductions), while premium reductions only occur at risk tier thresholds.

In summary, cybersecurity ratings are the currently the most important input to cyberinsurance calculations as a financial risk-transfer mechanism that allows organizations to exchange uncertain and potentially severe cyber losses for a predictable premium [54-58]. By transferring a portion of cyber risk to insurers, organizations reduce financial volatility and protect against events that cannot be fully prevented through cybersecurity controls alone. Cyberinsurance complements cybersecurity investments by absorbing residual financial risk, while organizations retain operational risk. The next two sections address the use of ratings as input to cybersecurity investments decisions meant to manage operational risk.

6.2 Cybersecurity Ratings as Input to the Gordon–Loeb Model

While cyberinsurance premiums reflect worst-case survivability, cybersecurity investments reduce expected operational loss. The Gordon–Loeb Model for Cybersecurity Investments (GL) answers the core question: How much should an organization rationally spend to protect an individual information asset from cyber risk in a given timer period? [64-69] GL treats cybersecurity spending as an economic decision, not a technical one. The key idea is that spending more on security reduces risk, however each additional dollar produces increasingly smaller benefits (diminishing returns), so this model identifies the upper bound on rational security spending (for a single asset in a single time period).

Given these GL inputs:

- ***L* — Loss**
  - monetary loss if the asset is compromised
  - includes costs for response, downtime, liability, brand damage, etc.
- ***v* — Vulnerability**
  - probability an asset is compromised without additional protection ***(This is where cybersecurity ratings fit: they provide an estimate of baseline vulnerability)***
- ***z* — Security investment**

  - amount spent on protecting an asset
- ***s(z)* — Security effectiveness**
  - function for how investment reduces vulnerability
  - assumes diminishing returns

GL Expected Loss Formulation:

Before investing: $Expected\ Loss_0 = EL_0 = L \times v$

After investing z: $Expected\ Loss(z) = L \times v \times (1 - s(z))$

Total organizational cost: $Total\ Cost(z) = z + L \times v \times (1 - s(z))$

The goal is to choose z that minimizes total cost**.** Under very general assumptions, GL proved the optimal security investment should be less than ~37% of the expected loss.

$$Formally: z^* \leq 0.37 \times (L \times v)$$

This result holds regardless of the exact form of the security effectiveness function and the specific compensating controls implemented. GL is not saying that an ideal cybersecurity budget is 37% of expected loss but rather that spending more than 37% is economically irrational for a single asset under standard assumptions, a reasonableness bound, not a budget target.

The GL model has important limitations:

- It assumes a single asset
- It assumes a single time period (does not consider the time value of money for losses or benefits over multiple years)

Because of these limitations, GL will undervalue:

- Program-level security investments over multiple years that have cost/benefit breakeven points.
- Compensating controls that protect multiple assets
- Required/mandatory investments for legal or organizational compliance
- Investments to reduce catastrophic risk for rare high-profile unforeseen events that have no precedent (Black Swan events).[5]

[5] A Black Swan event is an unpredictable, rare occurrence with extreme, widespread impact that are often inappropriately rationalized with after-the-fact hindsight that was not expressed prior to the event. Coined by Nassim Nicholas Taleb, these events defy typical forecasting, examples include the 2008 financial crisis, COVID-19 pandemic, and the 9/11 terrorist attacks [see below]. Computability for consequential rare events is problematic due to the nature of scientific calculations with small probabilities and human individual/collective psychological bias to estimating unforeseen, unprecedented events. Because these events cannot be predicted, conventional risk management techniques often fail. Strategies for handling them include building

Cybersecurity ratings contribute to GL as currently the most accurate way to quantitatively assess cybersecurity risk and thus are currently the best information source to estimate v (baseline vulnerability) – to provide accurate information about organizational exposure to cybersecurity losses before an investment decision.

6.3 Cybersecurity Ratings as Input to Net Present Value (NPV) Analysis

Net Present Value (NPV) analysis is the standard financial method used to decide whether an investment creates economic value over time - if benefits outweigh costs for the desired time period of investment. When comparing benefits to costs monetary loss estimates are modeled using domain-specific loss assumptions and time value of money assumptions.

Cybersecurity ratings support NPV analysis by validating reductions in incident likelihood with improved ratings and the durability of those reductions longitudinally over time. Incident likelihood, as reflected in cybersecurity ratings, represents the relative probability that an organization will experience certain types of cyber incidents based on externally observable empirical cybersecurity metrics, not a prediction of specific events or their impact.

NPV analysis sums all future cash inflows and outflows, discounted back to the present:

$$\text{NPV} = \sum_{t=0}^{T} \frac{\text{Benefits}_t - \text{Costs}_t}{(1+r)^t}$$

Where:

- t = time period
- r = discount rate (time value of money + risk)
- T = analysis horizon

Decision rule:

- *NPV > 0* → investment creates value → *economically justified*
- *NPV < 0* → investment destroys value → *do not invest*

Breakeven Point – time when investment *cumulative benefits > cumulative costs*

- If multiyear time period, Year when *NPV > 0*

NPV analysis lets you compare:

- doing nothing vs making a specific monetary investment

"robust" systems, such as portfolio diversification, maintaining high liquidity, and avoiding overly complex models that may give the psychological impression of forecasting certainty. [N.N. Taleb, "The Black Swan: The Impact of the Highly Improbable (second ed.)," New York: Random House Publishing Group, 2010.]

- comparing between Option A vs Option B (or other options C, D, E, etc…)
- large program portfolio investments vs incremental isolated investments

For cybersecurity, benefits are manifested by reduced expected losses (fewer incident responses), reduced downtime, avoided regulatory penalties and litigation, lower cyber-insurance premiums, and other operational efficiencies. For cybersecurity, costs are typically manifested by upfront capital costs, ongoing operating expenses, and training and staffing.

Cybersecurity ratings are a quantitative source for calculating benefits from reduced expected losses based on strategic investments designed to reduce risk (as reflected in ratings improvement) and thus resulting in a reduction in expected losses.

Given the domain-specific context of a cybersecurity investment, the NPV formula can evolve to the following specific formulation:

$$NPV = -C_0 + \sum_{t=1}^{T} \frac{\Delta EL_t}{(1+r)^t}$$

Where:

- C_0= upfront investment
- ΔEL_t= annual expected loss reduction *{loss reduction = positive benefit}*
  $\Delta EL = L \times (v_{before} - v_{after})$
  {assuming v_{after} will decrease in value for positive benefits}
- r= discount rate
- T= time horizon

ΔEL calculation in more detail:

Step A — Estimate baseline expected loss

$$EL_{\text{before}} = L \times v_{\text{before}}$$

Where:

- L= loss if compromised (estimate from business documentation)
- v= baseline vulnerability / likelihood *(supported by cybersecurity ratings)*

Step B — Estimate post-investment expected loss

$$EL_{\text{after}} = L \times v_{\text{after}}$$

Where:

- v_{after}reflects reduced likelihood due to compensating security controls

Step C — Compute *ΔEL annual benefit used in NPV*

$$\Delta EL = L \times (v_{\text{before}} - v_{\text{after}})$$

In summary, NPV tells you whether a cybersecurity investment creates value over time; GL tells you when cybersecurity spending for a single asset becomes economically unreasonable. While NPV is a valuation model that can be used to compare cybersecurity investments, GL model is a spending constraint upper bound on how much is reasonable to spend on cybersecurity for a single asset at a given time instance. *Cybersecurity ratings improvement calculations are a quantitative source for calculating benefits from reduced expected losses.*

6.4 Cybersecurity Ratings and the Integration of NPV and GL Analysis

NPV and GL analysis should be used together in cybersecurity investment decisions. NPV first evaluates whether investments create value over time by comparing discounted benefits to costs. GL can then provide an economically grounded upper bound on rational security spending based on expected loss, ensuring investments are not disproportionate to risk.

Cybersecurity ratings provide an externally observable proxy for baseline vulnerability and a continuous signal for whether risk-reduction assumptions are actually materializing. Cybersecurity ratings perform these two specific functions:

1. *Inform baseline vulnerability estimates (v) used in GL*
2. *Validate (or falsify) NPV benefit assumptions over time*

In both NPV and GL, the hardest variable to estimate is: How likely is a cyber incident *before* we invest? This is difficult because incidents are rare, data is sparse, internal assessments are subjective, and self-reporting is biased. Cybersecurity ratings help by providing external, independent, empirical, quantitative data.

Specifically, for NPV – cybersecurity ratings introduce empirical discipline. NPV models assume benefits exist, benefits persist, and benefits scale with time. Without validation, NPV becomes a spreadsheet exercise vulnerable to optimism bias. Cybersecurity ratings can serve as a validation check. Since NPV depends on benefits over time, measuring cybersecurity ratings over time can check whether improvement persists over time. The implications are:

- if cybersecurity ratings confirm that benefits are stable over time then benefits can be capitalized
- if cybersecurity ratings show reversion (up and down movement) then benefits
    - benefits have either been overstated, and/or
    - benefit period duration needs to be shortened (meaning benefits will not last for currently stated horizon), and/or
    - discount rate should be increased (so future benefits should be discounted more heavily)

*Specifically for GL – cybersecurity ratings can answer the question - Is the selected variable **v** value used in the analysis accurate?* This can be validated by cybersecurity

ratings measurements after the investment – did the investment increase/decrease/not-change security posture as reflected in cybersecurity ratings after the investment? If the cybersecurity ratings do not improve (or worsen) then the variable ***v*** value used in the GL formulas was inaccurate and spending is not economically justified under these assumptions. Thus, cybersecurity ratings support GL veto conditions.

Thus, cybersecurity ratings enable NPV and GL analysis to function together by providing an empirical basis for estimating baseline vulnerability and validating risk-reduction assumptions over time. Ratings inform the likelihood component of expected loss calculations, supporting GL reasonableness bounds on spending for a given level of expected loss, making it an effective guardrail against over-investment, while sustained improvements in ratings provide evidence that assumed reductions in incident likelihood used in NPV models are real and durable. By serving as an independent, continuous signal of exposure, cybersecurity ratings impose discipline on cyber investment decisions and prevent overreliance on speculative or unvalidated benefit assumptions.

Overall, NPV evaluates whether an investment creates value over time by comparing discounted benefits to costs, enabling portfolio-level decisions. Used together, GL constrains the decision space while NPV selects the best investments within it. Figure 10 summarizes comparison between NPV and GL analysis.

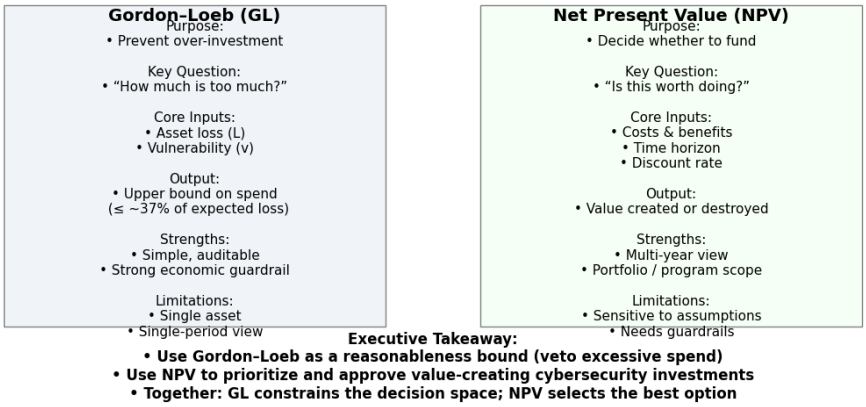

Figure 10. Cybersecurity Investment Decisions: Gordon-Loeb vs Net Present Value

Here is a practical outline for using NPV and GL together in a workflow:

1. ***Use ratings to estimate baseline v in GL***
2. Using GL compute EL_0 and apply GL upper bound
    - reject or rescope single asset investments which violate GL upper bound
3. Use desired cybersecurity ratings improvement to justify ΔEL in NPV
    - if NPV value has volatility or reversion, adjust NPV discount rates and/or benefit assumptions

4. Approve only if overall portfolio NPV calculation is positive for the given time period and individual asset investments are within GL upper bounds.
5. Monitor cybersecurity ratings post-investment to continuously reaffirm assumptions since conditions may change dynamically

Note that NPV and GL can disagree with a viable positive investment in this particular instance:

- GL single-asset bound violated
- *NPV > 0*
- The NPV scope is over a portfolio of multiple assets, not asset-specific.
- One single asset within the portfolio may provide a GL single-asset bound violation in which case this single asset may either be removed from the portfolio investment *or remain in portfolio investment if integral to operations*

Lastly, and perhaps most importantly, cybersecurity ratings are powerful not only because they can validate successful analysis NPV and GL assumptions, but also because they can be a falsification control to detect incorrect NPV and GL financial decisions. If a cybersecurity investment claims to reduce risk but improvement is not reflected in cybersecurity ratings, then the NPV benefit assumption is invalid and any amount of money spent under the GL upper-constraint is wasted.

In summary, cybersecurity ratings inform baseline vulnerability estimates and validate risk-reduction assumptions, allowing the two models to work together as complementary tools in disciplined cyber-risk governance.

7 Summary and Conclusions

In this paper we introduced cybersecurity ratings as an empirical technique to quantitatively measure an organization's security posture. We demonstrated the use of cybersecurity ratings in real-world environments: (1) to baseline the cybersecurity posture of two large domain-specific national infrastructures – all the hospitals in the U.S. and all the hospitals in Brazil – and (2) to consider possible financial cybersecurity investments by calculating return-on-investment (ROI) results for financial decision-making comparisons.

A cybersecurity baseline of domain-specific national infrastructure provides a consistent, high-level view of the overall security posture and exposure of a critical sector, enabling policymakers to identify systemic weaknesses, compare risk across industries, monitor trends over time, and prioritize investments and interventions. While not capturing all threats, such a baseline offers an essential situational awareness tool for understanding national-level cyber risk and resilience.

Using cybersecurity ratings to calculate quantitative ROI allows organizations to translate improvements in security posture into defensible estimates of risk reduction,

enabling cyber investments to be evaluated and compared using the same financial criteria as other capital decisions. By providing an external signal of changes in incident likelihood, ratings help validate ROI assumptions, prioritize investments with the greatest economic impact, and support disciplined, auditable allocation of limited financial resources.

This real-world practical experience leveraging cybersecurity ratings led to theoretical discussion illustrating how cybersecurity ratings are currently being incorporated in cyberinsurance calculations. We presented how cybersecurity ratings provide insurers with an independent, continuously updated indicator of an organization's relative cyber risk, particularly its likelihood of experiencing certain types of cyber incidents. Insurers use ratings to screen applicants, tier risk, inform underwriting decisions, and shape pricing, coverage limits, and exclusions. While ratings do not predict breach severity or guarantee outcomes, they enable scalable risk differentiation and ongoing monitoring, making cyber-insurance underwriting feasible.

We presented how cybersecurity ratings are important inputs to the Gordon-Loeb cybersecurity investment model and the Net Present Value multiyear portfolio decision-making process. In the Gordon–Loeb model, ratings inform the baseline likelihood of cyber incidents, allowing expected loss to be estimated and economically reasonable spending bounds to be set. In Net Present Value analysis, sustained improvements in ratings validate assumptions about reductions in incident likelihood and the durability of those reductions over time, supporting credible estimation of expected loss reduction (ΔEL). Together, ratings supply empirical evidence that allows Gordon–Loeb to constrain spending and NPV to evaluate value creation, ensuring cybersecurity investments are both economically reasonable and financially justified.

We conclude that cybersecurity ratings are a relatively new technique that will increasingly grow in importance due to the multi-dimensional utility we have shared in this paper. Cybersecurity ratings provide an independent, comparable signal of cyber risk and improvement over time, enabling informed prioritization, financial decision-making, insurance underwriting, and governance in domain-specific contexts that would otherwise not be possible with any accuracy.

Acknowledgments

We acknowledge and thank our WEIS 2026 peer reviewers for sharing their time and expertise to provide specific constructive feedback to us which we have incorporated to improve this paper.

References

[1] M.A. Goedeker, Building A Cyber Fusion Center with Advanced Threat Hunting and Intelligence: Mastering Threat Intelligence, Hunting, & XDR in a Few Weeks, independently published, 2024.

[2] D. Nathans, Designing and Building A Security Operations Center, Syngress, 2015.
[3] A. Basta, N. Basta, W. Anwar, and M.I. Essar, Open-Source Security Operations Center (SOC), Wiley, 2025.
[4] SPLUNK, 10 Essential Capabilities of a Modern SOC, Report 22-18111-Splunk-10-Essential-Capabilities-of-a-Modern-SOC-109, 2022.
[5] S.M. Adeel, AI-Driven Transformation of Security Operations Centers (SOCs), independently published, 2024.
[6] R. Blair, Aligning Security Operations with the MITRE ATT&CK Framework, Birmingham UK: Packt Publishing, 2023.
[7] J. Muniz, The Modern Security Operations Center, Addison-Wesley, 2021.
[8] K.L. McLaughlin, Cybersecurity Operations and Fusion Centers. CRC Press, 2024.
[9] K. DeValk and N. Elmqvist, "Riverside: A Design Study on Visualization for Situation Awareness in Cybersecurity," Information Visualization Journal 23(1) 2024.
[10] INFOSEC Research Council, 2005 Hard Problem List. Nov 2005. https://www.infosec-research.org/docs_public/20051130-IRC-HPL-FINAL.pdf
[11] M. Blaze, "Afterword" within B. Schneier, 1996. Applied Cryptography 2nd edition, 1996.
[12] National Institute of Standards and Technology (NIST), Measurement Guide for Information Security: Volume 1 – Identifying and Selecting Measures, NIST SP 800-55 Vol 1, Jan 2024.
[13] N. Bartol, B. Bates, K. Mercedes Goertzel, and T. Winograd, Measuring Cybersecurity and Information Assurance. State-of-the-Art Report (SOAR). DoD Information Assurance SOAR Technology Analysis Center (IATAC), May 8, 2009.
[14] S.M. Bellovin, On the Brittleness of Software and the Infeasibility of Security Metrics. IEEE Security & Privacy. 4(4) 96 July/August 2006. DOI:10.1109/MSP.2006.101.
[15] D.J. Bodeau, R.D. Graubart, R.M. McQuaid, and J. Woodill, Cyber Resiliency Metrics, Measures of Effectiveness, and Scoring. MITRE Technical Report, Release Case Number 18-2579, 2018.
[16] D. Chapin and S. Akridge, How Can Security Be Measured? Info Systems Control J, Vol 2 2005.
[17] J-H. Cho, P. Hurley and S. Xu, Metrics and Measurements of Trustworthy Systems. IEEE Military Comm Conf (MILCOM) 2016.
[18] L.F. DeKoven, A. Randall, A. Marian, G. Akiwate, A. Blume, L.K. Saul, A. Schulman, G.M. Voelker, and S. Savage, Measuring Security Practices. Comm of ACM. Sept 65(9) 2022. DOI:10.1145/3547133.
[19] D. Flater, Bad Security Metrics – Part 1: Problems. IEEE IT Professional, Jan/Feb 2018.
[20] D. Flater, Bad Security Metrics – Part 2: Solutions. IEEE IT Professional, Mar/Apr 2018.
[21] F. Innerhofer–Oberperfler and R. Breu. Potential Rating Indicators for Cyberinsurance: An Exploratory Qualitative Study. Workshop on the Economics of Information Security (WEIS), 2009.
[22] W. Jansen, Directions in Security Metrics Research, NIST Internal Report 7564, Apr 2009.
[23] G. Jelen, SSE-CMM Security Metrics, NIST and CSSPAB Workshop Washington DC, 2000.
[24] R. Khudhair and A. Ahmed, Overview of Security Metrics, Software Engineering, 4(4) 2016. DOI: 10.11648/j.se.20160404.11
[25] P. Manadhata and J.M. Wing, "An Attack Surface Metric," CMUCS-05-155 Carnegie Mellon University. 2005.
[26] M. Pendleton, R. Garcia-Lebron, J-H Cho, and S. Xu, A Survey on Systems Security Metrics. ACM Computing Surveys, 49(4) 1-35. Dec 2016. https://doi.org/10.1145/3005714.
[27] S.L. Pfleeger. 2009. Useful Cybersecurity Metrics. IEEE IT Professional. May/June 2009.
[28] S.L. Pfleeger and R.K. Cunningham. 2010. Why Measuring Security is Hard. IEEE Security & Privacy. July/August 2010.
[29] A.S. Pope, R. Morning, D.R. Tauritz, and A. D. Kent, Automated Design of Network Security Metrics. ACM Genetic and Evolutionary Computation Conference (GECCO) 2018.
[30] W. H. Sanders. Quantitative Security Metrics: Unattainable Holy Grail or a Vital Breakthrough within Our Reach? IEEE Security & Privacy. 12(2), Mar/Apr 2014. DOI:10.1109/MSP.2014.31.
[31] R.M. Savola, Towards a Taxonomy for Information Security Metrics. Intl Conf. on Software Engineering Advances (ICSEA), 2007.
[32] S. Schechter, Quantitatively Differentiating System Security. Workshop on the Economics of Information Security (WEIS), 2002.
[33] D. Snyder, L.A. Mayer, G. Weichenberg, D.C. Tarraf, B. Fox, M. Hura, S. Genc, and J.W. Welburn, Measuring Cybersecurity and Cyber Resiliency, RAND 2020. DOI:10.7249/RR2703.
[34] S. Stolfo, S.M. Bellovin, and D. Evans, 2011. Measuring Security. IEEE Security & Privacy, 9(3) May/Jun 2011. DOI:10.1109/MSP.2011.56.
[35] M. Torgerson, Security Metrics, 12th Intl Command and Control Research and Technology Symp, 2007.

[36] R.B. Vaughn, A. Siraj, and D.A. Dampier, Information Security System Rating and Ranking, Cross-Talk: The Journal of Defense Software Engineering, May 2002.
[37] R.B. Vaughn, A. Siraj, and R. Henning, Information Assurance Measures and Metrics—State of Practice and Proposed Taxonomy. 36th Hawaii Intl Conf on System Sciences (HICSS-36), January 2003.
[38] G.O. M. Yee, Designing Good Security Metrics. IEEE Annual Intl Computer Software and Applications Conference (COMPSAC) 2019.
[39] J. Zalewski, S. Drager, W. McKeever and A.J. Kornecki, Measuring Security: A Challenge for the Generation. Federated Conference on Computer Science and Information Systems, 2014. DOI: 10.15439/2014F490
[40] National Institute of Standards and Technology (NIST), National Vulnerability Database (NVD)/Known Exploited Vulnerabilities, <https://nvd.nist.gov/general/news/cisa-exploit-catalog>
[41] MITRE, CVE Program Mission. <https://www.cve.org/>
[42] W. Yurcik, S. North, R. O'Kane, O.S. Saydjari, F.R. Miranda, R.S. Avelino, and G. Pluta, Measurability: Toward Integrating Metrics into Ratings for Scalable Proactive Cybersecurity Management, Intl Conf on Emerging Security Information, Systems and Technologies (SECURWARE), 2025.
[43] National Institute of Standards and Technology (NIST), Guide to Enterprise Patch Management Planning: Preventive Maintenance for Technology, NIST SP 800-40 Rev. 4 Apr 2022.
[44] M.S. Ragheb, W. Elmedany, and M.S. Sharif, The Effectiveness of DKIM and SPF in Strengthening Email Security, 10th Intl Conf on Future Internet of Things and Cloud (FiCloud), 2023.
[45] S.J. Choi and M.E. Johnson, The Relationship Between Cybersecurity Ratings and the Risk of Hospital Data Breaches, J of the American Med Informatics Assoc. 28(10) 2021.
[46] D. Zängerle and D. Schiereck, Modelling and Predicting Enterprise-level Cyber Risks in the Context of Sparse Data Availability, Geneva Papers on Risk and Insurance - Issues and Practice, Palgrave Macmillan, The Geneva Association, 48(2), pp 434-462, Apr 2023.
[47] W. Yurcik and A. Schick, Change Healthcare – Perspective and Lessons from the Nationwide Pharmacy Supply Chain Failure, Cybersecurity in Healthcare, 2025. DOI:10.1109/ACSACW69556.2025.00021.
[48] Centers for Medicare & Medicaid Services (CMS) National Health Expenditures (NHE) 2024 Fact Sheet https://www.cms.gov/data-research/statistics-trends-and-reports/national-health-expenditure-data/nhe-fact-sheet
[49] J. Paim, C. Travassos, C. Almeida, L. Bahia and J. Macinko, The Brazilian Health System: History, Advances, and Challenges, Lancet, 377(9779), 2011. DOI:10.1016/S0140-6736(11)60054-8
[50] F. Ortega and A. Pele, Brazil's Unified Health System: 35 years and Future Challenges, The Lancet Regional Health - Americas, Vol 28, Nov 2023. DOI:10.1016/j.lana.2023.100631
[51] Brazilian Institute of Geography and Statistics, Estimates of the Resident Population in Brazil and Federation Units with Reference Date of Jul 1 2025. https://www.ibge.gov.br/estatisticas/sociais/populacao/9103-estimativas-de-populacao.html
[52] Brazil Ministry of Health, National Registry of Health Facilities (CNES), Department of Informatics of the Unified Health System (DATASUS), 2025. <http://cnes.datasus.gov.br/>
[53] Health-Information Sharing and Analysis Center (Health-ISAC) and RANE, The Brazilian Critical Infrastructure Threat Landscape and Implications for Healthcare Organizations, Threat Intelligence Report, May 2025. https://health-isac.org
[54] W. Yurcik and D. Doss, A Market Solution to the Internet Security Market Failure, Workshop on Economics of Information Systems (WEIS), 2002. <https://www.cl.cam.ac.uk/archive/rja14/econws/53.pdf>
[55] J.P. Kesan, R.P. Majuca, and W. Yurcik, The Economic Case for Cyberinsurance, Illinois Law and Economics Working Papers Series Paper No. LE04-004, 2004. https://law.bepress.com/uiuclwps/art2/>
[56] J.P. Kesan, R.P. Majuca, and W. Yurcik, Cyber-insurance As A Market-Based Solution To The Problem of Cybersecurity: A Case Study, Workshop on Economics of Information Systems (WEIS) 2005. <https://tinyurl.com/mpjkxr8j>
[57] R. P Majuca, W. Yurcik, and J.P. Kesan, The Evolution of Cyberinsurance, arXiv, 2006. https://arxiv.org/abs/cs/0601020
[58] J.P. Kesan, R.P. Majuca, and W. Yurcik, Three Economic Arguments for Cyberinsurance Chapter 16 within A. Chander, L. Gelman, & M. J. Radin (editors) Securing Privacy in the Internet Age, Stanford University Press, pp. 345-366, 2008. <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=577862>
[59] K. Awiszus et al. Modeling and Pricing Cyber Insurance. European Actuarial Journal, 13, 2023. DOI:10.1007/s13385-023-00341-9
[60] A. Tsohou et al. Cyber Insurance: State of the Art, Trends and Future Directions. Int. J. Inf. Secur. 22, 2023. DOI:10.1007/s10207-023-00660-8

[61] M.M. Khalili, M. Liu, S. Romanosky, Embracing and Controlling Risk Dependency in Cyber-Insurance Policy Underwriting, Journal of Cybersecurity, 5(1) 2019. DOI:10.1093/cybsec/tyz010
[62] D. Arce, D. W. Woods, and R. Böhme, Economics of Incident Response Panels in Cyber Insurance. Comput. Secur. 140, 2024. DOI:10.1016/j.cose.2024.103742
[63] R. Böhme and G. Kataria, On the Limits of Cyber-Insurance. In: Fischer-Hübner, S., Furnell, S., Lambrinoudakis, C. (eds) Trust and Privacy in Digital Business. TrustBus 2006. Lecture Notes in Computer Science, vol 4083. Springer, Berlin, Heidelberg 2006. DOI:10.1007/11824633_4
[64] L.A. Gordon and M.P. Loeb, The Economics of Information Security Investment. ACM Trans. Information System Security, 5, 2002.
[65] L.A. Gordon and M.P. Loeb. The Economics of Information Security Investment. within Camp, L.J., Lewis, S. (eds) Economics of Information Security. Advances in Information Security, Vol 12. Springer, 2004. DOI:10.1007/1-4020-8090-5_9
[66] J. Willemson, On the Gordon & Loeb Model for Information Security Investment. Workshop on the Economics of Information Security (WEIS) 2006.
[67] Y. Baryshnikov, IT Security Investment and Gordon/Loeb's 1/e Rule, Workshop on the Economics of Information Security (WEIS) 2012.
[68] H.R.K. Skeoch, Expanding the Gordon-Loeb Model to Cyber-Insurance, Computers & Security. Vol 112, 2022.
[69] A. Ebel and D. Mitra, Economics and Optimal Investment Policies of Attackers and Defenders in Cybersecurity, Journal of Cybersecurity. 10(1) 2024. DOI:10.1093/cybsec/tyae019