Measuring Return on Risk Reduction

A Modern Approach to Return on Security Investment (ROSI)

March 6, 2026

Authors:

Caleb Stogner
Senior Manager, Technology and Data Risk Management
Capital One
caleb.stogner@capitalone.com

Laura Voicu
Cofounder, Chief Data Science Officer
Enterprise Risk Quantification Institute
laura.voicu@me.com

Table of Contents

Executive Summary

Cybersecurity leaders are increasingly expected to justify security investments with the same financial rigor applied to other business initiatives. Traditional Return on Security Investment (ROSI) calculations offer a quick way to estimate the cost-effectiveness of security controls, but they often overlook important factors such as uncertainty in cyber risk, diminishing returns from additional spending, and the time value of money. As a result, ROSI alone can lead to incomplete or misleading conclusions when evaluating significant cybersecurity investments.

This paper presents a practical framework that combines FAIR-based cyber risk quantification with established financial decision tools such as Net Present Value (NPV), Internal Rate of Return (IRR), and the Gordon–Loeb model. FAIR provides a structured method for translating cyber risk into financial terms, enabling organizations to estimate how much loss a control is expected to prevent. These quantified risk reductions can then be evaluated using standard capital allocation methods that account for investment timing, opportunity cost, and long-term value creation.

For CISOs and other senior cybersecurity leaders, this approach provides a more defensible way to communicate the economic value of security initiatives and prioritize investments. By integrating risk quantification with proven financial analysis techniques, organizations can move beyond simplistic ROSI calculations and make cybersecurity spending decisions that are more transparent, economically sound, and aligned with enterprise value.

Introduction

Figuring out how much to invest in cybersecurity is a persistent challenge for many organizations. Leaders often grapple with the question, “Are we spending too little or too much?” and "Which of these options gives me the most value for my money?" These are fundamental questions that Return on Investment (ROI) is designed to answer. ROI is frequently used to compare alternative investment strategies by measuring profitability in terms of gains or revenue enhancements from an investment. However, classic ROI does not work for security investments because the traditional equation applies only to investments with predictable cash flows.

Security investments fundamentally differ: they neither increase revenues directly nor provide immediate positive cash flows. Rather, security investments are about risk management that results in loss prevention and risk mitigation. As a result, the term Return on Security Investment (ROSI) is often used to emphasize that cybersecurity calculations measure loss avoidance rather than profit generation. Here, the notion of return is broad, as prevented losses do not constitute returns in the traditional financial sense. Unfortunately, traditional methods for measuring the effectiveness of security investments, ROSI, repeat mistakes that have been solved in other professional domains. By the traditional ROSI method, we are specifically referring to variations of the following formula:

$$ROSI = (ALE * M - C) / C$$

ROSI = Return on Security Investment
ALE = Annual Expected Loss
M = Mitigation Ratio between 0 - 100%
C = Cost of Control

While easy to calculate and communicate, ROSI overlooks crucial factors such as the uncertainty of future threats, the diminishing benefits as spending increases, and the fundamental idea that money available now is worth more than money available later. To make smarter, long-term cybersecurity decisions, organizations need tools that can address these gaps. This paper introduces more robust tools from financial decision science to address these gaps.

The Case for Evolving Beyond ROSI

Despite its shortcomings, ROSI remains popular because it offers a straightforward way to measure short-term, tactical security decisions. It clearly communicates immediate benefits, making it easy for decision-makers to justify spending. However, the simplicity it creates can introduce noise that must be calibrated to the decision at hand. The key gaps in traditional applications of ROSI include:

- They fail to account for the time value of money, effectively treating future and current benefits as equally valuable.

- They don't account for uncertainty arising from the rapidly evolving threat landscape, causing controls to lose effectiveness over time.
- They often encourage short-term thinking over long-term strategic planning.

For some types of decisions, the noise introduced is worth the time saved by applying the lighter, traditional method, especially for short-term or clearly defined tactical decisions such as targeted upgrades or narrowly scoped risk scenarios. The key is recognizing when ROSI fits the decision context and when it's necessary to complement it with more robust financial and risk analysis methods for strategic, long-term cybersecurity planning.

Introducing New Tools: NPV/IRR, Gordon Loeb Model

NPV/IRR

Net Present Value (NPV) and Internal Rate of Return (IRR) are financial metrics that evaluate the expected returns on both capital and operating expenditures over time. These two interrelated tools answer the fundamental question of how well investments perform against a given discount rate[1], or, in other words, how well investments perform against a stated target of performance level. The selected discount rate is a key assumption that should be determined in collaboration with the organization's finance function. Often, the organization's cost of capital[2] is selected as the hurdle rate for investment projects.

Both NPV and IRR use present-value factors in discounted cash flow analysis. Because organizations seek to sustain cash flows and fuel growth, spending patterns are typically compared against their ability to meet those objectives. Understanding opportunity costs (i.e., the value of choosing the best alternative investment) usually prevents organizations from seeking lower-growth opportunities. All other things being equal, if a cybersecurity investment would return 8% to the business but spending on new product features would return 20%, the best decision is obvious.

How might cybersecurity spending fit into this mode of decision-making, given that the security function is usually a cost center rather than a revenue generator? For practitioners of FAIR, the value of security is recognized as the reduction in expected future losses, or reduced annualized loss exposure (ALE). Said more simply, the primary benefit of security is risk reduction. But risk reduction is rarely free, and deciding *how much to spend* on it is the key question.

[1] In ROI analysis (including NPV and IRR), the discount rate is the assumed annual rate of return used to convert future costs and benefits into present-day dollars, reflecting the organization's time value of money and risk tolerance.

[2] The cost of capital is the organization's blended cost of funding (debt ain expected future lossesthat a 15% growth opportunity, it is not difficult to decide where to allocatend equity) and represents the minimum acceptable return required to preserve firm value

Producing a discounted cash flow ("DCF") analysis first requires the analyst to identify the expected cash outflows (i.e., both capital investments and operating expenditures) and the anticipated benefits (i.e., reduction in ALE), and then allocate them to their respective time periods (usually years) over a given duration. Several examples will be provided towards the end of this paper to help illustrate how to arrange these analyses.

Back to IRR and NPV as decision tools, each of these financial metrics can be derived from a completed DCF analysis. How they differ in their purpose and use is this:

- NPV will return a single monetary amount that reveals the project's overall benefits, expressed in today's dollars. If an NPV is positive, the project is considered profitable, as the benefits outweigh the costs (and outperform the selected discount rate).
- IRR will return a single rate expressed as a percentage, analogous to the return percentages you see with bonds or other securities (e.g., this bond returns 6% over 5 years). Mathematically, IRR is the discount rate that makes the NPV equal to $0.

NPV and IRR are effectively two sides of the same coin. They both illustrate the value of an investment (benefits minus costs) over a defined period (e.g., three years). And they both reflect the time value of money for the organization.

When comparing multiple alternatives (e.g., different security investments competing for a limited budget, evaluating different risk scenarios), NPV will often be a more suitable tool for the job of evaluating the spend and risk reduction for the following reasons:

- NPV does not require that the same discount rate be used for all alternatives; as a result, it is more modular for more complex analysis scenarios:
    - The timeframes of the competing alternatives vary in length and cost structure
    - The scenarios addressed by the competing alternatives have little or no overlap or have varying degrees of confidence

That being said, if you are measuring mutually exclusive mitigation alternatives for one combination of risk scenarios for a given stakeholder (i.e., should we pursue mitigation A or mitigation B), IRR can be a quicker time-to-value than NPV because:

- The fundamental risk scenarios being measured are, for the most part, identical or at least matched in duration and confidence levels.
- IRR analyses do not require a discount rate and can be performed if that data point has not yet been provided.

As for required inputs, the only difference between NPV and IRR is that the former requires a discount rate. Therefore, choosing one over the other is often straightforward: when a discount rate is available, both NPV and IRR are calculated for analysis. The discussion above aims to guide the practitioner in which metric to emphasize with stakeholders based on the decisions to be supported.

Gordon Loeb Model

The Gordon-Loeb Model, first introduced in the influential 2002 paper *"The Economics of Information Security"* by Lawrence A. Gordon and Martin P. Loeb, is widely recognized as a foundational framework for deciding how much organizations should invest in cybersecurity. Its primary insight is that optimal cybersecurity spending should typically be a fraction of the expected loss from a breach, rather than equal to or greater than the potential loss.

The model's most cited result is that security spending should rarely exceed about 37% of the expected loss. While this figure might seem arbitrary, it is the result of testing the model across a wide range of assumptions about how security investments reduce risk and observing that the optimal spending point never rises above this level. Gordon and Loeb explored many ways security investments might reduce risk, from aggressive scenarios (where each dollar is highly effective) to conservative ones (where each dollar yields only small improvements). Across all these variations, the model consistently showed that the optimal spending level was always less than 37% of the expected loss. In other words, no matter the assumptions, it is almost never economically rational to spend more than about a third of what you stand to lose. In practical terms, the 37% threshold should be viewed as an upper bound rather than a precise prescription, rather like a benchmark that helps ensure discipline and avoid overspending.

The basic idea is one of diminishing returns: the first dollars invested in security tend to have the greatest impact, while subsequent dollars deliver progressively smaller reductions in risk. After a certain point, additional spending has little effect on expected loss and may even become economically inefficient.

It is also important to note that, like all models, the Gordon-Loeb model comes with certain assumptions and limitations. One such limitation is that the model requires a predictable, smooth relationship between investment and risk reduction. In other words, as you spend more, risk declines in a steady and continuous way. This matters because the model's optimization logic depends on that steady relationship. In practice, however, cybersecurity spending is rarely smooth. Investments are "lumpy": a tool purchase, a staff member hire, or a process change, and each of these steps reduces risk in jumps rather than along a continuous curve. For this reason, the Gordon–Loeb model works best as a guidepost rather than a precise calculator.

Another limitation is that the model treats expected loss as a single fixed value and assumes a predictable effect of spending, while real-world risk is inherently probabilistic. This, however, is where FAIR complements Gordon–Loeb to overcome this limitation. FAIR models risk as distributions and scenarios, capturing uncertainty and producing a realistic range of potential losses, while Gordon–Loeb provides a simple economic guardrail by capping spend at roughly a third of expected loss.

Bringing the two approaches together balances detail with discipline: FAIR quantifies risk in probabilistic terms, and Gordon–Loeb reminds decision-makers that it is rarely efficient to spend

more than a fraction of the losses they seek to prevent. The model's role is less about calculating an exact "right number" and more about reminding decision-makers that it is rarely justified to spend more than a fraction of the loss you are trying to avoid.

Framework (The How)

Optimizing cybersecurity spending requires a repeatable process that translates risk into consumable inputs to financial decisions. Our proposed framework integrates FAIR for risk modeling with economic evaluation methods (ROSI, NPV, IRR, and the Gordon–Loeb cap) in a transparent, auditable, and adaptable manner to evolving threats. It is designed to (i) scope and quantify risk rigorously, (ii) map controls to measurable changes in that risk, and (iii) express outcomes in the language of capital allocation and enterprise value.

1. Model Risk Scenarios with FAIR

The framework begins by defining decision-relevant risk scenarios and quantifying them using FAIR. The unit of analysis is a scenario that specifies the asset or service at risk, the threat community, the relevant control environment, and loss types. FAIR's decomposition forces clarity: frequency drivers (contact rates, action rates, resistance strength versus threat capability) are modeled separately from magnitude drivers (direct response costs, replacement and productivity effects, and second-order impacts such as regulatory exposure or dependent losses).

Uncertainty is expressed explicitly through probability distributions on the FAIR parameters rather than point estimates. This accounts for both random variation in how events unfold and uncertainty arising from gaps or limits in available information. When internal data are sparse, FAIR relies on structured expert-judgment protocols, calibrated elicitation, and external reference data, with all assumptions recorded for auditability. Correlations across scenarios, assets, or loss drivers are documented to avoid double-counting and to reflect shared causes (for example, a widespread vulnerability or a critical vendor).

The output of this step is an enterprise-level loss distribution and scenario-level Annualized Loss Exposure (ALE) estimates that are traceable to their inputs, versioned over time, and suitable for downstream financial analysis.

2. Model Mitigation Spending and Control Effects

With scenario baselines established, the next task is to model how specific controls change FAIR factors and what they cost in practice. Controls are treated as interventions that modify the frequency or magnitude of loss events, or both. The framework requires a defensible mapping from each proposed control to the FAIR factors it affects, along with any assumptions explaining the causal pathways (for instance, improved detection shortening dwell time, or process improvements

reducing secondary losses). For further reading, refer to the FAIR Institute's materials on the FAIR Control Analytics Model (FAIR-CAM™).

Control implementation costs should be modeled comprehensively as lifecycle cash flows: acquisition, implementation, and integration; recurring operating costs; induced costs, such as performance overhead or false-positive handling; decommissioning or transition costs at the end of life; and organizational changes, such as training, policy change, or process redesign. Effectiveness is time-dependent: adoption curves, deployment lags, learning effects, and potential control decay should be captured to avoid assuming immediate and permanent efficacy. Overlaps and interactions among controls should be addressed explicitly to reflect diminishing returns and to prevent double-counting when multiple interventions target the same FAIR factors.

Where empirical measurement is feasible, our framework encourages pre-defined evaluation plans (e.g., staged rollouts, counterfactual baselines, or comparable groups) to reduce reliance on untested assumptions. When measurement is constrained, the analysis documents uncertainty bands and identifies leading indicators that must be monitored post-deployment to validate or update the model.

The output of this step is a set of alternative mitigation portfolios, each represented by (a) a revised loss distribution for the affected scenarios and (b) a schedule of cash flows capturing the total cost of ownership and operational impacts.

3. Apply Economic Metrics for Decision Quality

Economic evaluation then translates risk changes and costs into decision-ready metrics that align with enterprise finance.

ROSI remains useful for small, well-scoped initiatives with clear, near-term effects and short payback horizons. In this framework, ROSI is a screening tool for tactical decisions where uncertainty is limited and delivery risk is low.

NPV is the primary metric for multi-year or strategic initiatives. Benefits are modeled as reductions in expected loss net of induced costs; discounting reflects the organization's cost of capital and includes a risk premium appropriate to the uncertainty of security outcomes over time. The framework avoids "double-discounting" by being explicit about whether uncertainty is represented via risk-adjusted discount rates or via probability distributions in the cash flows. When the finance function publishes a WACC, the analysis anchors to it and documents any premiums applied for project-specific uncertainty. IRR provides a hurdle-rate comparison, useful for ranking alternatives when capital is rationed. The Gordon–Loeb model provides an economic upper bound. For each scenario, the proposed spend is compared to a cap derived from expected loss to guard against economically inefficient over-investment. At the portfolio level, this cap should be treated as a constraint rather than a target.

4. Uncertainty, Sensitivity, and the Value of Information

Uncertainty is unavoidable in cyber risk analysis, so it needs to be addressed directly. The framework recommends treating sensitivity analysis as a core activity: test how much results change when key assumptions shift. For example, what happens if attack frequency is twice as high as expected, or if control effectiveness is lower than assumed? Stress tests should also explore extreme but plausible cases, such as simultaneous failures or concentrated losses from a single vulnerability.

When decisions rely on poorly understood parameters, it can be worth estimating the value of gathering better information through targeted measurement, pilot projects, or external sources. This avoids wasting effort on refining inputs that won't affect the choice, while focusing resources where better data could materially improve the decision.

Applying The Methods with Guardrails

Selecting Between ROSI and an NPV/IRR Analysis

Having covered the methods and approach to deriving the NPV and IRR for a given security mitigation project, when should an organization exercise these methods over the simpler return on security investment (ROSI)? The practicality of the simple ROSI formula popularized within the industry has led to widespread adoption and makes it adequate for certain types of decisions. Some screening criteria that may indicate that ROSI would be sufficient for the use case are presented below, as a non-exhaustive list:

- The in-scope security mitigation project:
    - has a payback period somewhere between 12 and 24 months; or
    - has an implementation period of less than 12 months; or
    - needs a quick go/no-go decision for a small-scale implementation.

In situations where longer timeframes are necessary to capture the spending and benefits of the mitigation project, discounting metrics such as NPV and IRR will more holistically reflect the risk against the organization's required rates of return.

Applying Guardrails

There are two main guardrails within discounted cash flow metrics that can be adjusted to restrain overspending and capture uncertainty. The first is inherent within the discount rate selected. For investments with a greater risk of failure (i.e., where costs or benefits are less certain), organizations may apply a "risk premium" in addition to their standard discount rate. This standard practice accounts for the risk tolerance in the use of money. The uncertainty introduced by the FAIR analysis as an input to the discounted cash flow analysis could, in part, be addressed by applying some

premium expressed as a percentage in addition to the agreed discount rate, so that, absent other data, a concrete return from cash inflows will be favored over an expected loss reduction of the same amount. The second guardrail is the Gordon-Loeb constraint of 37% for limiting spend towards risk reduction.

Measurement Examples

Bank Ransomware Scenario

To further illustrate these methods in action, consider a risk scenario in which ransomware is the attack outcome, and spear-phishing is the initial attack method, targeting a bank. The full composition of the potential loss event is described as follows:

> **Threat**: Financially motivated cybercriminal groups using Ransomware-as-a-Service (RaaS). These well-resourced gangs specialize in ransomware extortion, encrypting data and demanding payment.
>
> **Asset**: The bank's critical assets, including customer personal and financial data (e.g., account details, Social Security numbers) and transaction systems vital for business continuity.
>
> **Method**: The attacker gains initial access via a targeted spear-phishing email with a malicious attachment or link. While other vectors (e.g., software vulnerabilities) exist, phishing is the primary focus for this example.
>
> **Effect**: The attack causes:

- **Business Interruption**: Downtime in online/mobile banking, leading to lost revenue and operational backlogs.
- **Direct Financial Loss**: Ransom payments (if paid), incident response costs (e.g., forensics, IT remediation), and staffing expenses.
- **Regulatory and Legal Penalties**: Fines for leaked PII, breach notification costs, and legal fees from lawsuits.
- **Reputation Damage**: Eroded customer trust, impacting long-term confidence and share value.
- **Opportunity Costs**: Missed business opportunities (e.g., new loans, trades) and increased insurance premiums.

The bank is evaluating a proposed investment in e-mail filtering software that's expected to prevent malicious e-mail from being delivered to employees' inboxes. The cost of this project totals $1 million, with $700K expected in the initial year for implementation costs and $100K each year thereafter for subscription costs under a 3-year contract.

Without this e-mail filtering solution in place, the organization's anticipated annualized loss exposure (ALE) is $850K. The e-mail filtering solution is expected to reduce the ALE of this risk scenario by 66.7% to $283K, primarily by reducing the organization's susceptibility to this style of attack (i.e., fewer malicious links or attachments reach user inboxes to be clicked). This is a net annualized reduction of $567K in risk, but is that worth the $1 million price tag over the 3-year contract?

- When looking at the first-year economics and comparing the $700K implementation against the Year 1 Risk Reduction, the traditional ROSI looks appealing:
    - ($850K - $700K) / $700K = 21% ROSI
- When extrapolating the scenarios out to the full 3-year contract's duration, the ROSI looks even better:
    - ($1.7 million - $1.0 million) / $1.0 million = 70% ROSI

Are either of these measurements, on their own, enough to inform a final decision of whether to pursue the security mitigation? If you have other alternatives with similar ROSIs to compare to, they might be. However, they do not ensure that the return on the investment exceeds the company's cost of capital, as measured by a finance-approved discount rate. For a more thorough analysis, see the following NPV, IRR, and Gordon-Loeb analyses, using the same fact patterns to enable improved decision-making across more use cases.

First, assume this organization's finance department has informed the risk analyst that the discount rate to use for the DCF is 15%. The project's economics and ALE should be summarized by year:

Example #1: Bank Ransomware Scenario – Undiscounted				
	Up-Front	**Year 1**	**Year 2**	**Year 3**
Average Annualized Loss Exposure ($)	-	850,000	850,000	850,000
Average Annualized Loss Exposure ($) After Investment (Residual Risk)	-	283,333	283,333	283,333
Risk Reduction ($)	**-**	**$566,667**	**$566,667**	**$566,667**
Costs of Mitigation	(700,000)	(100,000)	(100,000)	(100,000)
Net Risk Mitigation Benefits	**($700,000)**	**$466,667**	**$466,667**	**$466,667**

When applying present value discounting to the individual year's net risk mitigation benefits above at 15%, an NPV of $365.5K. How is this calculation performed? As a refresher, let's take each year's bottom-line number (the "Net Risk Mitigation Benefits") and apply it through the following formula:

$$PV = amount * (1 + discountrate)^{-year}$$

$-\$700{,}000 * (1 + 0.15)^{0} = -\$700{,}000$
$\$466{,}667 * (1 + 0.15)^{1} = \$405{,}797$

$466,667 * (1 + 0.15)^2 = $352,867
$466,667 * (1 + 0.15)^3 = $306,841

The present value over three years, including the up-front investment, totals an NPV of $365,505. This calculation is possible when a defensible discount rate is available, say, from the finance team, but what if one is not? The IRR metric can be calculated without a discount rate, as it imputes the break-even discount rate available to set the NPV to $0. This can be done manually through testing or with purpose-built methods in Python or formulas in Excel to perform the interpolation automatically. In this case, recognize that the IRR for the example above is 44.6%.

> ***How do we know this?*** *Apart from utilizing available formulas, here is a test:*
>
> *$-700,000 * (1 + 0.446)^{-0} = -700,000*
> *$466,667 * (1 + 0.446)^{-1} = 322,730*
> *$466,667 * (1 + 0.446)^{-2} = 223,188*
> *$466,667 * (1 + 0.446)^{-3} = 154,348*
>
> *The sums from the new discounting analysis at 44.6% differ only due to rounding, with an NPV of approximately $300. The closer we pushed the discount rate to the actual IRR of 44.631201%, the closer that would approach $0.*

How would the Gordon-Loeb model evaluate this mitigation spending? With the 37% threshold (i.e., security spending should not exceed $0.37 for each dollar of risk mitigated), the comparison suggests that the organization should spend no more than about $315K per year to address this risk. At first glance, that appears to rule out the project entirely, since the upfront $700K implementation cost is more than double the cap. But this is where Gordon–Loeb shows both its limitation and its value. Security investments are rarely spread evenly over time. Instead, they are "lumpy," with large initial outlays followed by smaller recurring costs. If the cap is applied strictly on an annual basis, many practical projects will look inefficient. In practice, the Gordon–Loeb model can be applied either annually or over the full project horizon. On a strict annual view, the $700K upfront implementation cost clearly exceeds the $315K yearly cap. But if we compare the three-year totals, the cap rises to about $943K ($850K * 3 * .37), while actual spend is $1M, so the project comes in only slightly above the threshold.

The real value of Gordon–Loeb in this example is not to deliver a simple yes/no, but to flag that the proposed mitigation is heavily front-loaded relative to the expected annual loss reduction. This prompts decision-makers to ask whether the spend can be structured differently or whether there are alternative controls that achieve similar benefits at a lower upfront cost. Gordon–Loeb works best here as a guardrail: it highlights potential overspend, while tools like NPV and IRR capture the full economics of cost and benefit flows over time.

So, in the example of the Bank facing a Ransomware scenario, the return calculations are as follows:

Metric	Measurement
Traditional ROSI	70.0%
Net Present Value (NPV)	$366K
Internal Rate of Return (IRR)	44.6%

Additional Examples

Not all organizations will have the same impacts in a ransomware scenario and may have unique financing or capital structures that would alter the scenario above. For that reason, two additional example organizations are included to show how the economics might shift across different industries, as opposed to the banking-focused scenario.

Professional Services Organization Ransomware Example #2

Consider a consulting organization that primarily provides professional services to its clients and has relatively low capital. When assessing a similar ransomware scenario to the one above, the organization's cyber risk analysis team estimated the ALE at $500K, given a lower expected impact. Furthermore, they believe that the e-mail filtering control will be more effective than the previous scenario, reducing the likelihood by 80% rather than the bank's estimate of a 66.7% reduction. For purposes of this example, let's also assume this organization has selected a discount rate of 15%. The economics of this organization's project are summarized below:

Example #2: Professional Services Ransomware Scenario – Undiscounted				
	Up-Front	Year 1	Year 2	Year 3
Average Annualized Loss Exposure ($)	-	500,000	500,000	500,000
Average Annualized Loss Exposure ($) After Investment (Residual Risk)	-	100,000	100,000	100,000
Risk Reduction ($)	**-**	**$400,000**	**$400,000**	**$400,000**
Costs of Mitigation	(700,000)	(100,000)	(100,000)	(100,000)
Net Risk Mitigation Benefits	**($700,000)**	**$300,000**	**$300,000**	**$300,000**

At a 15% discount rate, the NPV of this project is negative, at -$15K. The IRR of this project resolves to 13.7%, which explains the negative NPV. Anytime the IRR is below an organization's selected discount rate, the NPV will be negative. Higher IRRs than discount rates will result in positive NPVs. With a baseline ALE of $500K, the Gordon–Loeb cap would be about $185K per year. The project costs of $700K upfront plus $100K annually are well above this threshold, reinforcing the negative NPV result.

Here, the model highlights the same concern: the organization is likely overspending relative to the risk reduction, and alternatives should be explored.

Manufacturing Organization Ransomware Example #3

Consider a manufacturing organization that is heavily focused on capital discipline, preferring to defer spending to prioritize manufacturing more goods. When assessing a similar ransomware scenario as previously laid out, their CRQ department arrived at an ALE estimate of $1.2 million, given the heightened sensitivity to business continuity disruptions and the larger supply chain impacts if manufacturing is halted. The control effectiveness estimates of 66.7% that the banking organization landed on were also considered appropriate for this organization.

One change this organization managed was to negotiate the payment terms with the vendor, so that the costs of implementation are more spread out than in the other two examples. Initial implementation costs will be only $400K, with $200K paid each year thereafter to reach the same $1 million total; in other words, spending is not as heavily front-loaded as that of the other examples.

For purposes of this example, let's also assume this organization has selected a discount rate of 15%. The economics of this organization's project are summarized below:

Example #3: Manufacturing Org Ransomware Scenario – Undiscounted				
	Up-Front	**Year 1**	**Year 2**	**Year 3**
Average Annualized Loss Exposure ($)	-	1,200,000	1,200,000	1,200,000
Average Annualized Loss Exposure ($) After Investment (Residual Risk)	-	400,000	400,000	400,000
Risk Reduction ($)	**-**	**$800,000**	**$800,000**	**$800,000**
Costs of Mitigation	(400,000)	(200,000)	(200,000)	(200,000)
Net Risk Mitigation Benefits	**($400,000)**	**$600,000**	**$600,000**	**$600,000**

At a 15% discount rate, the project's NPV is $970K. The IRR of this project resolves to 139%. With a baseline ALE of $1.2M, the Gordon–Loeb cap is about $444K per year. The project's staggered costs ($400K upfront, $200K annually) come closer to that level, and over the three-year horizon, the total spend of $1M stays well within the model's multi-year cap of about $1.3M. This aligns with the strong NPV and IRR results.

Example Recap

Each of the organizations in the examples above is evaluating very similar facts: (1) A ransomware risk scenario originating from a spear-phishing attack vector and (2) the possibility of committing to an implementation project with ongoing subscription costs for an email filtering control to mitigate that risk. What changes around those two sets are facts are:

- How likely and/or impactful the given risk scenario will be against the target organization
- The target organization's individual capital needs and the anticipated effectiveness of the control

The table below summarizes the three cases:

Example Scenario	Bank	Pro Services Firm	Manufacturer
Inputs			
Cost (Y0)	$700,000	$700,000	$400,000
Cost (Recurring)	$100,000	$100,000	$200,000
Duration (years)	3	3	3
ALE	$850,000	$500,000	$1,200,000
Risk Reduction %	66.67%	80.00%	66.67%
Metrics			
Discount Rate	15%	15%	15%
NPV	$365,506	-$15,032	$969,936
Payback Period (years)	1.8	2.5	1.2
Traditional ROSI	70%	20%	140%
Gordon-Loeb Spending Cap	$943,500	$555,000	$1,332,000
Full Spending	$1,000,000	$1,000,000	$1,000,000
Security IRR	**44.6%**	**13.7%**	**139.0%**

Conclusion

Cybersecurity should no longer be seen merely as a defensive expense or a compliance checkbox, but as a strategic investment that directly supports business resilience and long-term value. The challenge lies in making those investment decisions with the same rigor applied to other areas of corporate finance.

By combining the structured risk quantification of FAIR with established financial tools such as NPV, IRR, and the Gordon–Loeb Model, organizations can move beyond intuition-driven or fear-driven budgeting. This integrated approach enables leaders to translate technical risk into economic terms, balance short-term and long-term considerations, and ensure that spending levels remain both effective and economically rational.

Adopting such a framework allows organizations to make cybersecurity investment decisions that are transparent, defensible, and aligned with enterprise value objectives, elevating security from a cost center to a disciplined component of strategic financial management.

Glossary of Terms

ALE (Annualized Loss Exposure): In FAIR, ALE represents the expected annual loss from a specific risk scenario, calculated as the product of *Loss Event Frequency* (how often losses occur) and *Loss Magnitude* (the expected cost per event). It provides a common financial unit for comparing and prioritizing risk reduction options.

DCF (Discounted Cash Flow): A valuation method that estimates the present value of future cash flows by applying a discount rate. DCF is used to assess whether an investment's expected returns justify its cost, accounting for the time value of money.

NPV (Net Present Value): The sum of discounted future cash inflows minus the initial investment cost. A positive NPV indicates that an investment is expected to generate value beyond its cost, after adjusting for risk and time.

IRR (Internal Rate of Return): The discount rate at which an investment's NPV equals zero. IRR represents the expected rate of return generated by the project or control investment over its lifecycle.

Discount Rate: The rate used to convert future cash flows or losses into present value terms. It reflects the organization's cost of capital (e.g., WACC) and the perceived risk of the investment or control measure.

ROI (Return on Investment): A general financial metric expressing the ratio of net benefits (gains minus costs) to total costs. ROI provides a high-level measure of investment efficiency or profitability.

ROSI (Return on Security Investment): (*Traditional definition)* A cost-effectiveness measure for security where benefits derive from loss avoidance, often operationalized as the reduction in expected loss relative to control cost. *(From this paper)* A measure of the economic value created by a cybersecurity control or investment, expressed as the reduction in expected loss relative to the cost of the mitigation. Here, the measurement is derived from FAIR-based estimates of risk reduction (e.g., decreases in Annualized Loss Exposure) and is evaluated using standard financial decision tools such as discounted cash flow analysis, Net Present Value (NPV), and Internal Rate of Return (IRR) to determine whether the investment creates economic value for the organization.

References

- Gordon, Lawrence A.; Loeb, Martin P. (2004). "Economics of Information Security Investment". In Camp, L. Jean; Lewis, Stephen (eds.). *Economics of Information Security*. Advances in Information Security. Vol. 12. Boston, MA: Springer.
- Caleb Stogner (2023). “Redefining ROSI in Risk Assessment: A Practical Guide for Risk Analysts”, FAIR Institute
- Laura Voicu (2025). “Bringing Financial Discipline to Cyber-Risk Decisions – A Practitioner’s Field Guide”, FAIR Institute
- 2025 State of Cyber Risk Management Report, FAIR Institute