FROM THREAT INTELLIGENCE TO ECONOMIC METRICS: LLM-SUPPORTED CYBER RISKOPS FOR CONTINUOUS SV/ROSI ASSESSMENT

CZU: 004.056:005.342
DOI: https://doi.org/10.53486/tids2026.27

GHENADIE BELINSCHI
Academy of Economic Studies of Moldova (ASEM), Doctoral School, Doctoral Programme 521.03 "Economics and Management (Business and Professional Practice)", Republic of Moldova
e-mail: ghenadie.belinschi@ase.md
ORCID ID: 0009-0009-3361-9890;

IRINA DOROGAIA
Academy of Economic Studies of Moldova (ASEM), Department of Management and Entrepreneurship, Republic of Moldova
e-mail: dorogaia.irina.ion@ase.md
ORCID ID: 0000-0003-4625-8616

Abstract. This paper examines the gap between rapidly evolving cyber threats and the relatively static economic justification of cybersecurity innovations in organizations. Building on a risk-based view of information management and on the concept of Cyber RiskOps as a continuous risk-management cycle, we propose an LLM-supported pipeline that translates Threat Intelligence (TI) into auditable updates of scenario parameters used in economic metrics. The pipeline combines an information system passport—a structured register of information assets, business processes, dependencies, and controls—a scenario library, and external/internal signals. The large language model (LLM) is used strictly as an extraction and normalization component, not as an autonomous decision-making mechanism. It transforms unstructured TI into structured attributes, including actor, TTPs, affected technologies, exploitation maturity, and business-impact drivers, and maps them to scenarios through transparent scoring rules and human verification. The updated probabilities and loss magnitudes serve as inputs for the Saved Value of Information (SV) metric, defined as the avoided losses related to confidentiality, integrity, and availability, and for its linkage with Return on Security Investment (ROSI) when comparing security controls under an as-a-service subscription model. An illustrative SME case shows how faster parameter updates can change investment decisions: a package of controls that appears unprofitable under outdated assumptions may become ROSI-positive once TI-driven probability shifts are incorporated in a timely manner. The paper discusses governance requirements, including prompt and data control, audit trails, and segregation of duties, as well as limitations of the approach, such as uncertainty in loss estimation and LLM-specific risks, including hallucinations and prompt injection. The proposed approach supports doctoral research on improving information management through cybersecurity innovations by operationalizing continuous, evidence-based measurement of the economic value of security.

Keywords: LLM; threat intelligence; cyber risk quantification; ROSI; cybersecurity innovations; information management.

JEL Classification: M15; D81; O33

INTRODUCTION

Digital transformation makes information a key factor of production and a critical intangible asset for organizations. This is consistent with the information management perspective, which treats information as a managed resource (Rowley and Farrow, 2000). However, the economic effect of cybersecurity is largely indirect: rather than generating new revenue, security measures primarily

prevent or reduce losses caused by incidents, regulatory sanctions, downtime, and reputational damage. This complicates the investment case and requires decision-makers to rely on risk-based loss-prevention metrics, such as Return on Security Investment (ROSI) (Sonnenreich, Albanese and Stout, 2005), as well as other economic models of security investment (Gordon and Loeb, 2002; Beissel, 2016).

At the same time, the threat landscape is highly dynamic. Current threat reports document the growth of identity-based attacks, business email compromise (BEC), and malware-free intrusions, while generative AI reduces the cost of social engineering and increases the personalization and scalability of attacks (CrowdStrike, 2025; Check Point Research, 2025). In this context, static or infrequently updated risk registers produce outdated estimates of probabilities and loss magnitudes. As a result, SV/ROSI calculations may diverge from actual exposure, leading either to underinvestment in relevant controls or to overinvestment in measures that address an outdated threat profile.

In the doctoral research devoted to improving information management through cybersecurity innovations, we propose a measurable link between: (i) the "information system passport" as a structured representation of information assets, processes, and dependencies; (ii) scenario-based risk assessment; and (iii) the Saved Value of Information (SV) metric, understood as the value of information preserved through the avoidance of losses related to confidentiality, integrity, and availability. This triadic approach supports a managerial perspective in which information management and cybersecurity are integrated and assessed through measurable effects, rather than merely through compliance checklists.

This paper further develops the proposed approach by introducing a continuous updating mechanism inspired by Cyber RiskOps, which treats cyber-risk management as an operational cycle linking strategy, threat intelligence, execution of controls, and measurable outcomes (Castro, 2025; Leirvik, 2023). We argue that large language models (LLMs) can reduce analytical latency within this cycle, but only if their use is limited to auditable data-extraction tasks and they are embedded in the broader management process.

The study is guided by two research questions: RQ1: How can threat intelligence be translated into updates of scenario parameters for SV/ROSI calculations in an auditable and reproducible manner? RQ2: Under what conditions can faster parameter updates change the economic decision regarding cybersecurity innovations for organizations operating under subscription-based models?

The contribution of this paper is threefold. First, we propose an LLM-supported Cyber RiskOps pipeline that links unstructured Threat Intelligence (TI) with a scenario library and an information system passport. Second, we formalize the relationship between TI-driven updates and SV/ROSI calculation under uncertainty and governance constraints. Third, we illustrate the approach through a demonstration SME case study and discuss the limitations and control measures required for the responsible use of LLMs in risk analytics.

MAIN CONTENT

1. LITERATURE REVIEW

The evaluation of cybersecurity investments has a long history of combining technical risk analysis with economic decision-making. Classical approaches operationalize risk as the product of probability and impact and use derived metrics such as Single Loss Expectancy (SLE) and Annualized Loss Expectancy (ALE). In the corporate context, these metrics are often linked to ROI-type analysis through ROSI, which expresses net savings, understood as avoided losses, relative to the cost of a

security control (Sonnenreich, Albanese and Stout, 2005). The Gordon–Loeb model further emphasizes that optimal security expenditure is typically only a fraction of expected losses and depends on the vulnerability level of the information set being protected (Gordon and Loeb, 2002).

Despite their usefulness, economic metrics face practical limitations: loss magnitudes are uncertain, some impacts are intangible, and probability estimates are sensitive to both internal conditions and external threat dynamics (Beissel, 2016). As a result, many organizations still communicate cyber risks through qualitative heat maps. Modern approaches to cyber risk quantification (CRQ), such as FAIR (Factor Analysis of Information Risk), seek to improve consistency by decomposing risk into factors such as threat event frequency, vulnerability, and loss magnitude, while encouraging calibration, interval estimates, and Monte Carlo simulation (FAIR Institute, 2025). Frameworks and standards such as the NIST Cybersecurity Framework 2.0 and ISO/IEC 27001/27005 provide a process-oriented structure for identifying and managing cybersecurity risks and for linking security controls to desired outcomes (NIST, 2024; ISO/IEC, 2022a; ISO/IEC, 2022b).

From an operational perspective, the challenge lies not only in how metrics are calculated, but also in how their relevance is maintained over time. The concept of Cyber RiskOps proposes an operating model in which cyber risks are managed continuously across functional boundaries: strategy defines risk appetite and objectives, operations implement controls, and measurement continuously informs prioritization (Castro, 2025). Similar ideas are also found in the practical literature on cyber-risk management, which emphasizes continuous monitoring and integration with business processes (Leirvik, 2023).

Threat intelligence is one of the key external inputs that should update risk assumptions. However, in security operations centers (SOCs), TI is often consumed primarily as indicators and tactics, rather than translated into economic parameters. A significant gap lies in the absence of a standardized and auditable process for transforming unstructured TI narratives into the structured attributes required to update scenario probabilities and loss magnitudes.

Large language models (LLMs) offer the potential to reduce the manual workload associated with processing TI. They are effective in summarization, entity and relationship extraction, and mapping text to predefined taxonomies. At the same time, their use introduces new risks, including hallucinations, prompt injection, and uncontrolled leakage of sensitive data if models are deployed without appropriate governance. Professional communities increasingly emphasize that AI can both amplify threats and support defensive capabilities, which makes responsible implementation and control mechanisms essential (Check Point Research, 2025; NIST, 2025).

Thus, the research gap addressed by this paper lies in designing a governance-aware pipeline that uses LLMs as constrained extraction tools to integrate TI into continuous cyber-risk quantification and to support SV/ROSI-based decision-making in information management.

2. METHODOLOGY

The proposed approach is positioned as a measurement layer within the Cyber RiskOps cycle. Its purpose is to update scenario parameters—namely, probability and loss magnitude—on the basis of new threat-intelligence signals and to recalculate SV/ROSI for alternative cybersecurity innovations. The pipeline assumes that the organization already maintains a minimum set of governance artifacts typical of information-management and security programs: an inventory of information assets and systems, a control catalog, and a risk register. We structure these artifacts into three key components.

(1) Information System Passport (ISP). The ISP is a structured description of the organization's information-system landscape: critical business processes, information assets (datasets), supporting applications and infrastructure, owners, third parties, and dependencies on controls. Unlike a general IT inventory, the ISP is specifically designed to link risk to value: each asset is tagged with confidentiality, integrity, and availability (CIA) requirements and with business-impact drivers, such as acceptable downtime, regulatory exposure, and revenue dependency.

(2) Scenario Library. The organization defines a set of recurring cyber-risk scenarios relevant to its context. Each scenario k includes a brief description, a mapping to the affected ISP elements, a baseline probability p_k, a baseline loss magnitude L_k, and key control levers that can reduce p_k and/or L_k. Scenarios may be aligned with common threat categories, such as disruption caused by ransomware, business email compromise, identity takeover, cloud misconfiguration, or third-party compromise.

(3) Signal Streams. The pipeline uses both external and internal signals. External signals include TI reports, notifications, vendor bulletins, and industry alerts. Internal signals include incidents, near misses, audit results, and control-effectiveness metrics, such as MFA coverage, phishing click-through rate, and patch deployment delays.

For each scenario k, the expected losses at time t is defined as:

$$EL_k(t) = p_k(t) \times L_k(t) \quad (1)$$

where p_k(t) denotes the probability that scenario k materializes over the selected period, typically one year, and L_k(t) denotes the corresponding loss magnitude expressed in monetary terms. This definition is consistent with standard cyber-risk quantification practice and allows the SV metric to be interpreted as avoided expected losses.

LLM-supported extraction. In this context, the LLM is introduced as a text-to-structure component. The model is not allowed to directly assign risk values. Instead, it produces a structured extraction record for each TI item, including fields such as: (i) affected technologies and vendors; (ii) attack vectors and tactics/techniques, with mapping to MITRE ATT&CK where appropriate; (iii) sectoral relevance; (iv) exploitation maturity, such as theoretical, proof-of-concept, or observed exploitation; (v) potential business impacts and preconditions; and (vi) recommended mitigation measures. Each extraction record includes evidence pointers, such as text fragments and links, to ensure auditability.

Figure 1 summarizes the proposed pipeline and its main data flows.



Figure 1. LLM-supported Cyber RiskOps pipeline for the continuous updating of SV/ROSI.
Source: *author's own conceptualization.*

Matching and scoring. A deterministic matching function links each extraction record to one or more scenarios in the library according to explicit rules. For example, if TI indicates increased exploitation of cloud identity services and realistic voice impersonation, the matching function raises the signal score for scenarios related to identity compromise and BEC. The resulting score is then translated into scenario-parameter updates through bounded adjustments, for example:

$$p_k(t+1) = \text{clip}(p_k(t) \cdot (1 + \alpha \cdot S_k), p_{min}, p_{max}) \quad (2)$$

$$L_k(t+1) = \text{clip}(L_k(t) \cdot (1 + \beta \cdot I_k), L_{min}, L_{max}) \quad (3)$$

where S_k denotes the normalized signal score for probability, and I_k denotes the signal score for loss drivers. The parameters α and β are calibration coefficients selected by the organization. The clip() operator imposes predefined bounds and prevents overreaction to a single source. Importantly, final parameter updates must be approved by a human reviewer, such as the risk owner or an analyst, before they take effect in the risk register.

Economic calculation. The Saved Value of Information (SV) for a portfolio of controls c over a given period is calculated as the avoided expected losses across scenarios, comparing the baseline state (without the additional controls) with the post-implementation state (after the controls have been implemented):

$$SV_c = \Sigma_k [EL_k^{base} - EL_k^c] \quad (4)$$

where EL_k^base is calculated using the baseline parameters, and EL_k^c is calculated using the parameters after applying the portfolio of controls c. If the controls affect CIA dimensions differently, L_k can be decomposed into confidentiality, integrity, and availability components and then aggregated in a way that reflects the organization's priorities.

ROSI is used to normalize SV by the value of the control portfolio C_c:

$$ROSI_c = (SV_c - C_c) / C_c \quad (5)$$

This formulation is consistent with the ROSI literature and supports comparison across alternative control portfolios (Sonnenreich, Albanese and Stout, 2005). In subscription-based environments, C_c typically represents annual recurring costs, such as licences and managed services, whileSV_c is interpreted as annual avoided losses.

Accounting for uncertainty. Because both p_k and L_k are uncertain, the pipeline supports interval values rather than point estimates. An extraction record may include a qualitative confidence level, such as high, medium, or low, which is then converted into narrower or wider intervals. Organizations can perform sensitivity analysis by recalculating SV/ROSI across plausible parameter ranges or, when sufficient data are available, apply FAIR-style Monte Carlo simulation (FAIR Institute, 2025).

Governance and auditability. The use of LLMs in risk analytics requires controls similar to those applied to other high-impact decision-support systems. We propose the following minimum safeguards: (i) retrieval-only prompts that require the model to cite evidence fragments; (ii) separation of duties between TI ingestion, parameter approval, and the final investment decision; (iii) logging of prompts, model outputs, and reviewer decisions; (iv) data minimization, including a prohibition on sending sensitive internal logs to external models, and the use of secure environments; and (v) periodic validation against actual incidents and control-performance metrics in order to recalibrate α and β. The result is an auditable chain: TI source → structured extraction → scoring rule → parameter update → SV/ROSI calculation → management decision.

3. CASE STUDY

To demonstrate the impact of continuous parameter updates on the investment decision, we use a simplified SME case study. The organization is a professional services firm with approximately 25

employees, operating in a hybrid work model and relying on Microsoft 365, a cloud-based CRM, and online banking for daily operations. Following the dissertation approach, the firm maintains an Information System Passport (ISP) that identifies critical processes, including service delivery to clients, invoicing, and cash-flow management, as well as information assets, including client data, financial documents, and user credentials, together with their CIA requirements.

The scenario library includes three annual scenarios: (S1) ransomware-induced disruption of workstations and the file-sharing environment; (S2) business email compromise (BEC) and voice phishing (vishing) leading to fraudulent payments; and (S3) data leakage caused by SaaS misconfiguration or credential compromise. The baseline estimates for the next 12 months, prepared in Q1, are presented in Table 1 as the "outdated baseline (Q1)".

The cybersecurity innovation under consideration is an identity and email protection package that includes mandatory MFA with conditional access, enhanced payment-approval procedures, user training focused on AI-enabled social engineering, and a managed email-security service. The estimated annual recurring cost of the package is C = 1{,}600 EUR under a subscription model.

The expected control effect is modeled through scenario-specific reduction factors applied to the baseline probability and/or loss magnitude: the probability of ransomware disruption is reduced from 0.07 to 0.06; the probability of BEC is reduced by 50%, while the corresponding loss magnitude is reduced by 25%; and the probability of data leakage is reduced from 0.10 to 0.08, while the loss magnitude decreases from 15,000 EUR to 13,000 EUR.

With the outdated baseline, the total expected loss is EL_base = 4,690 EUR per year, while the expected loss after implementing the control package is EL_control = 3,080 EUR per year. This yields SV = 1,610 EUR per year and ROSI ≈ 0.6%. Under this assessment, the investment appears marginal and may be assigned a lower priority.

Next, the LLM-supported pipeline ingests new TI items describing an increase in AI-enabled impersonation and BEC campaigns targeting SMEs that use widely adopted cloud-based email platforms. After extraction, matching with scenario S2, and human review, the probability of BEC for the forthcoming period is updated from 0.12 to 0.22, forming the "TI-updated baseline." The remaining scenarios in this simplified example remain unchanged.

With this update, the total baseline expected loss increases to 5,890 EUR per year, while the expected loss after implementing the control package becomes 3,530 EUR per year, since the same relative control effect is applied to the updated probability. As a result, SV amounts to 2,360 EUR per year, and ROSI ≈ 47.5%. Therefore, continuous TI-based updating changes the investment decision: the same cybersecurity innovation becomes economically justified once the shift in the threat landscape is reflected in a timely manner.

Although the figures are illustrative, the case demonstrates the core mechanism: faster, evidence-based updates can change SV/ROSI results, re-prioritize control portfolios, and reduce the risk of resource misallocation caused by outdated assumptions. In practice, the approach can be extended by using interval estimates and by validating updates against internal incidents and control-effectiveness metrics.

Table 1. Scenario-Based Expected Losses under Outdated and TI-Updated Parameters: Demonstration Case.

Scenario	EL_base (outdated Q1), EUR	EL_control (outdated), EUR	EL_base (TI-outdated), EUR	EL_control (TI-outdated), EUR
S1: Ransomware-induced disruption	1,750	1,500	1,750	1,500
S2: BEC/vishing fraud	1,440	540	2,640	990
S3: SaaS data leakage	1,500	1,040	1,500	1,040
Total	**4,690**	**3,080**	**5,890**	**3,530**

Source: *authors' calculations.*

4. DISCUSSION AND LIMITATIONS

The proposed approach represents a pragmatic compromise between fully manual cyber-risk quantification and fully automated, opaque AI-based scoring. Its value lies not in replacing expert judgment, but in operationalizing the evidence flow from TI to economic metrics, which is often a bottleneck in organizations with limited analytical capacity. This is especially relevant in subscription-based environments, where cybersecurity innovations are purchased as a service and must be continuously validated against a changing exposure landscape.

However, several limitations must be acknowledged. First, loss assessment remains uncertain and context-dependent. Even with an information system passport, translating CIA requirements into monetary losses requires assumptions about downtime costs, recovery capabilities, contractual penalties, and reputational effects. Therefore, SV/ROSI results should be viewed as decision support under uncertainty rather than as precise forecasts.

Second, parameter updates may introduce volatility. Overreacting to a single report or a noisy signal can lead to fluctuations in prioritization. For this reason, the method relies on bounded updates via the clip function, calibration coefficients, and multi-source corroboration before substantial parameter changes are made. Organizations may also limit the update frequency, for example to a weekly cycle, in order to balance responsiveness and stability.

Third, LLM-specific risks require dedicated controls. Hallucinations may produce false mappings or “invent” facts if prompts are not sufficiently constrained. Prompt injection may occur when unverified TI content is passed to the model. Data-leakage risks arise when internal incident data or sensitive inventories are processed by external model providers. For these reasons, the use of LLMs should be limited to extraction from approved sources, supported by retrieval-based mechanisms and evidence citation, and embedded in an approval workflow with logging and segregation of duties. Emerging guidance on cybersecurity risk management for AI systems (NIST, 2025) may provide an additional governance reference point.

Finally, this paper demonstrates the approach using a simplified example. Full empirical validation would require longitudinal organizational data, comparison of decision outcomes with and without continuous updating, and assessment of analyst workload and error rates. These directions are planned as further work within the dissertation research programme.

CONCLUSIONS

This paper proposes an LLM-supported Cyber RiskOps pipeline for the continuous measurement of the economic effect of cybersecurity innovations. The central idea is to close the gap between dynamic threat intelligence (TI) and the relatively static risk parameters used in SV/ROSI calculations. The pipeline integrates three governance artifacts—the information system passport, the scenario library, and signal streams—while the LLM is used only for auditable extraction and normalization tasks.

By defining explicit matching and scoring rules and requiring human approval of parameter updates, the approach preserves transparency and supports an evidence-based audit trail from TI sources to updated scenario probabilities and loss magnitudes. The demonstration SME case study shows that timely TI-based updating can change the perceived economic attractiveness of a control portfolio: a package that appears marginal under outdated assumptions becomes ROSI-positive once the shift in the threat landscape is captured.

The proposed method supports the broader dissertation research programme on improving information management through cybersecurity innovations by providing a practical mechanism for the continuous measurement of value. Future work will focus on developing a dashboard prototype, integrating internal telemetry for improved calibration, and evaluating the approach on longitudinal datasets in order to quantify reductions in analytical latency and improvements in decision quality.

REFERENCES

Beissel, S. (2016) Cybersecurity Investments: Decision Support Under Economic Aspects of Safeguards and Breaches. Cham: Springer International Publishing. doi:10.1007/978-3-319-30460-1.

Castro, J.P. (2025) Cyber RiskOps: Bridging Strategy and Operations in Cybersecurity. ResearchGate. doi:10.13140/RG.2.2.36216.97282/1. Available at: https://www.researchgate.net/publication/388194428 [Accessed April 17, 2026].

Check Point Research (2025) AI Security Report 2025. Check Point Software Technologies. Available at: https://www.checkpoint.com/resources/items/report-ai-security-report-2025/ [Accessed April 17, 2026]

CrowdStrike (2025) European Threat Landscape Report 2025. CrowdStrike. Available at: https://www.crowdstrike.com/wp-content/uploads/2025/11/CrowdStrike-European-Threat-Landscape-Report-2025.pdf [Accessed 17 April 2026].

FAIR Institute (2025) Factor Analysis of Information Risk (FAIR) Standard v3.0 (January 2025). FAIR Institute. Available at: https://www.fairinstitute.org/hubfs/Standards%20Artifacts/Factor%20Analysis%20of%20Information%20Risk%20%28FAIR%29%20Standard%20v3.0%20%28January%202025%29.pdf [Accessed April 17, 2026].

Gordon, LA and Loeb, MP (2002) 'The economics of information security investment'. ACM Transactions on Information and System Security, 5(4), pp. 438–457. doi:10.1145/581271.581274.

ISO/IEC (2022a) ISO/IEC 27001:2022 Information security, cybersecurity and privacy protection — Information security management systems — Requirements. Geneva: International Organization for Standardization.

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

Leirvik, R. (2023) Understand, Manage, and Measure Cyber Risk: Practical Solutions for Creating a Sustainable Cyber Program. 2nd edn. Berkeley, CA: Apress.

NIST (2024) The NIST Cybersecurity Framework (CSF) 2.0. NIST CSWP 29. National Institute of Standards and Technology. Available at: https://doi.org/10.6028/NIST.CSWP.29 [Accessed April 17, 2026].

NIST (2025) Cybersecurity Framework Profile for Artificial Intelligence (AI) Profile: Preliminary Draft. NIST IR 8596 iprd. National Institute of Standards and Technology. Available at: https://nvlpubs.nist.gov/nistpubs/ir/2025/NIST.IR.8596.iprd.pdf [Accessed April 17, 2026].

Rowley, J. and Farrow, J. (2000) Organizing Knowledge: An Introduction to Managing Access to Information. 3rd edn. Aldershot: Gower.

Sonnenreich, W., Albanese, J. and Stout, B. (2005) 'Return on security investment (ROSI): A practical quantitative model'. In: Proceedings of the 3rd International Workshop on Security in Information Systems (WOSIS 2005). pp. 239–252. doi:10.5220/0002580202390252.