International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 11 | Nov - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Investing to Reduce Risk: Prioritizing Physical and Cybersecurity under Budget Constraints

Maximiliano Servilheira Sacaquini
Independent Researcher, Brazil
max.sacaquini@icloud.com
Corresponding author: Maximiliano Servilheira Sacaquini
Running title: Prioritizing Physical and Cybersecurity Investment
Submitted to: International Journal of Scientific Research in Engineering and Management (IJSREM)

***Abstract*—This paper develops an auditable procedure for prioritizing physical and cybersecurity investments when funding, staff capacity, and evidence are limited. Security budgets often separate physical and cyber protection even when both sustain the same service. A critical integrative review informs a finite-scenario binary portfolio model incorporating mandatory measures, prerequisites, lifecycle costs, and complementary controls. The working proposition is that ranking actions individually by value-to-cost ratio does not, in general, yield the feasible portfolio with the best worst-case assessed value. The objective maximizes the lowest assessed protection value across explicitly stated scenarios without interpreting stakeholder value points as expected monetary loss. An entirely hypothetical six-action instance is solved by exhaustive enumeration. Under the stated assumptions, the robust portfolio attains a minimum value of 12 points, compared with 9 for a specified ratio-based heuristic and 8 for the scenario-specific optimum. Sensitivity analysis identifies budget plateaus, staff bottlenecks, and infeasible recommendations caused by omitted prerequisites. These are mathematical properties of the constructed instance, not measured security improvements. The contribution is a transparent connection between prioritization, portfolio feasibility, uncertainty, and accountable acceptance of residual exposure.**

***Keywords*—security investment, portfolio selection, resource allocation, robust decision support, cyber-physical risk**

I. INTRODUCTION

An organization can afford a new surveillance platform yet lack the staff to maintain its permissions. It can strengthen network access while leaving the controller cabinet physically exposed. It can also purchase two controls that address the same weakness while postponing the capability on which both depend. In each case, the apparent investment decision concerns an individual product, but the consequential decision concerns a portfolio of protection capabilities.

The economic case for treating security expenditure as a decision problem is well established. Gordon and Loeb connect investment to vulnerability and potential loss, showing why the most vulnerable information set need not be the best investment target [1]. Gordon, Loeb, and Zhou subsequently relate cost–benefit reasoning to cybersecurity framework implementation [2]. These contributions do not supply a universal division of expenditure between guarding, physical access, connected equipment, digital protection, and operational recovery. That allocation requires an explicit account of the services at stake and the constraints governing delivery.

The challenge is particularly relevant where financial and organizational resources are constrained. Heidt, Gerlach, and Buxmann examine characteristics of small and medium-sized enterprises (SMEs) that affect information security investment [3]. Chidukwani, Zander, and Koutsakis review the uneven research coverage of security capabilities in smaller businesses [4]. The recent review by Rombaldo Junior, Becker, and Johnson further distinguishes evidence about SME cybersecurity from claims that circulate without empirical testing [5]. These studies motivate caution about assuming that organizations have specialist staff, stable cost records, or reliable loss estimates. They do not justify treating all smaller enterprises as a homogeneous population.

Physical and digital investments also interact. A door controller relies on hardware, power, permissions, communications, and procedures. A backup can preserve information without restoring the physical process that uses it. Integrated cyber-physical risk analysis, including the attention to assets and cascading effects developed by Kure et al. [6], provides a basis for asking how a portfolio sustains a protection service. The unit of analysis in this paper is therefore the service, rather than the department purchasing a control.

The research question is: how can an organization prioritize physical and cybersecurity investments under budget and staffing constraints when controls are interdependent and outcome estimates are uncertain? The operational scope is a commercial or institutional facility in which physical security, information technology, facilities, and external suppliers jointly support protected entry, monitoring, and continuity. No statutory size classification or single jurisdiction is assumed. The working proposition examined is that, when actions have prerequisites, compete for staff, or produce complementary value, an ordering of individual actions by value-to-cost ratio does not, in general, identify the feasible portfolio with the best worst-case assessed value. The proposition is examined through a formal portfolio model and an exactly solved hypothetical instance.

The paper makes three contributions. First, it separates mandatory service requirements from discretionary improvements, preventing a high score from compensating for an unacceptable safety or operational condition. Second, it formalizes portfolio selection with resource limits, prerequisites, and explicitly assessed complementarity. Third, it connects the numerical recommendation to a decision record that retains assumptions, rejected alternatives, and residual exposure. The optimization components are established

techniques; the contribution is their traceable application to joint physical and cyber investment decisions under limited evidence. No new general optimization algorithm is claimed.

Section II reviews the related literature. Section III describes the review method, the proposed decision procedure, the portfolio formulation, and the design of the computational illustration. Section IV reports the results of the illustration. Section V interprets these results and discusses governance, implementation, and limitations. Section VI concludes.

II. LITERATURE REVIEW

A. Economic Value and the Limits of Loss Estimates

The Gordon–Loeb model offers a disciplined economic interpretation of security spending [1]. Its findings depend on the specified breach probability functions, the potential loss, and the decision setting. A model-derived fraction of expected loss should not be converted into a universal budget rule for a facility combining personnel, physical assets, information, and life safety obligations. In particular, a monetary comparison cannot establish permission to omit a mandatory protective requirement.

The approach in Gordon et al. [2] shows how framework-based security management can be connected to economic choice. The present paper retains that demand for explicit trade-offs but uses a different decision representation when probabilities and losses cannot be estimated credibly. Stakeholder value points can expose preferences without pretending to measure avoided monetary loss. They introduce their own assumptions, especially comparability and cardinal interpretation, which must remain visible.

Security control selection has also been formulated as a portfolio problem. Sawik and Sawik use binary optimization and a cybersecurity value concept to select control portfolios in a supply-chain setting [7]. This establishes an important precedent for evaluating combinations rather than dividing a budget among independently ranked controls. The present formulation extends the application context to a joint physical and cyber service, with an explicit staffing limit and a decision record. It does not claim that binary portfolio selection or modelling control interaction is new.

B. Multiple Criteria, Dependencies, and Organizational Capacity

Sukumar et al. use multicriteria methods to assess cyber risks among UK small e-tailers [8]. Their inclusion of different risk themes is relevant to the need for a broader decision structure. Their empirical weights and rankings, however, are not transferred here: an e-tailing sample cannot determine the priorities of a different facility or security service. A priority ordering is also distinct from a feasible investment portfolio. Several individually attractive actions may not fit the same funding or staff envelope.

Cherdantseva et al. compare cyber risk assessment methods for supervisory control and data acquisition systems [9]. Their review highlights why the purpose and information demands of a method matter. Complex probabilistic representations can be valuable when inputs support them; they can also conceal uncertainty when probabilities are supplied primarily to make a calculation possible. The proposal here makes uncertain conditions explicit through a finite set of scenarios rather than assuming a distribution.

Dependencies can be technical, physical, or organizational. A monitoring integration may require protected administrative access. A door maintenance programme and a response exercise may provide an additional joint capability when staff learn to recognize and handle the maintained equipment's failure modes. Conversely, two measures may overlap substantially. Kure et al.'s integrated treatment of assets and cascading vulnerabilities [6] motivates recording these relationships before scores are aggregated.

Staff capacity is a separate resource. A supplier's installation price may exclude internal configuration, acceptance, training, and recurring review. The organizational constraints discussed by Heidt et al. [3] make that omission consequential. Kobis examines human-factor aspects of information security management in traditional and cloud settings [10]. The implication drawn here is limited but practical: purchasing a capability does not establish that people can operate it as intended.

C. Risk Guidance and Uncertainty

The NIST Cybersecurity Framework 2.0 provides organizational outcomes rather than a prescribed investment algorithm [11]. NIST IR 8286 Revision 1 connects cybersecurity risk information with enterprise risk management [12], while SP 800-30 Revision 1 supports structured assessment of threats, vulnerabilities, consequences, and uncertainty [13]. These documents motivate accountable decisions and explicit assumptions. They do not provide empirical effectiveness coefficients for the portfolio model below.

Robust optimization addresses decisions that must remain acceptable when model data change. Bertsimas and Sim demonstrate how conservatism can be managed in robust formulations [14]. The finite-scenario maximin model used here is simpler than their uncertainty-budget construction. It should be understood as a transparent decision convention for a declared scenario set, not as an implementation of their specific method or a probabilistic guarantee against future attacks.

III. METHODOLOGY

A. Research Design and Review Method

This is a critical integrative review with a proposed decision procedure and a constructed computational illustration. Snyder describes integrative reviews as a means of combining perspectives to develop conceptual frameworks [15]. The present work does not claim an exhaustive systematic search, a meta-analysis, or a field evaluation.

The starting bibliography, compiled on 19 October 2026, mapped security investment, SME constraints, risk assessment, and governance. Targeted searches were conducted on 2 November 2026 using Consensus and SciSpace. A Consensus query combined security investment, budget allocation, physical and cybersecurity, portfolio optimization, and interdependent controls. SciSpace was asked how organizations can prioritize physical security and cybersecurity investments while accounting for control dependencies, lifecycle costs, limited staff, and uncertainty. Publisher and institutional sources were then used to check selected references and the claims retained.

Selection was purposive. Included sources addressed economic choice, multicriteria assessment, cyber-physical dependencies, organizational capacity, optimization under uncertainty, or an institutional risk process. Duplicate versions were treated as one work. Some records were accessible through abstracts or publisher extracts rather than complete articles; those sources support only statements available at that

level. No clause-level claims are drawn from inaccessible licensed standards. Citation counts and journal metrics were not used as substitutes for methodological appraisal.

The synthesis organized material around the decision object, resources, non-negotiable requirements, interaction between actions, uncertainty, and decision ownership. These categories informed the proposed matrix and mathematical constraints. Literature findings, institutional recommendations, and the paper's own modelling conventions are distinguished throughout. The resulting model is not calibrated to any organization.

B. Service Definition and Admissible Alternatives

The process begins with a protection objective stated in operational terms. For example, a facility may require authorized entry, usable incident records, and a workable response during a local outage. The baseline describes what is already in place. Candidate actions must refer to incremental changes from that baseline, otherwise replacement expenditure and existing capability may be counted as new benefit.

The next step identifies mandatory measures. These may follow from established safety arrangements, applicable requirements, or explicitly approved minimum service conditions. Their status must be determined by an authorized person and recorded with its basis. A requirement is not made mandatory merely because a supplier describes it as essential. Equally, a mandatory action cannot be dropped because another action has a higher score. The integration of safety and security involves interactions that deserve separate scrutiny, as discussed by Kriaa et al. [16].

Candidates that introduce an unacceptable operating condition are redesigned or excluded before scoring. NIST's operational technology guidance recognizes the need to consider performance, reliability, and safety alongside cybersecurity [17]. Applying that perspective here means asking whether a proposed isolation measure interferes with a required emergency function, whether a change can be reversed, and who approves the resulting operating state. These questions are screening conditions, not extra points in an otherwise compensatory ranking.

TABLE I. PROPOSED INVESTMENT ASSESSMENT MATRIX AND ITS DECISION USE

Dimension	Evidence or information to record	Use in the decision
Service criticality	Protected service, users affected, duration and severity of failure	Define outcome criteria and mandatory limits
Incremental contribution	Baseline, failure mechanism addressed, expected functional change	Assess scenario-dependent value
Dependencies	Prerequisites, shared infrastructure, overlapping protection	Define feasible combinations and interactions
Lifecycle cost	Acquisition, configuration, support, renewal, retirement	Define comparable cost inputs and resource limits
Delivery capacity	Internal effort, specialist availability, acceptance and training	Apply a separate capacity constraint
Uncertainty	Evidence gaps, alternative assumptions, review triggers	Construct scenarios and sensitivity checks
Accountability	Sponsor, technical owner, acceptance authority, residual-risk owner	Approve implementation and exceptions

Table I is a proposed decision aid. Its dimensions are not multiplied into a risk score. Criticality shapes the outcome criteria and service limits; implementation feasibility becomes a constraint; uncertainty becomes an explicit modelling condition. This avoids counting the same concern repeatedly under several differently named scoring criteria.

C. Cost, Capacity, and Evidence

Costs must be compared over the same planning horizon. For a discretionary action i, a possible present-value input is

$$c_i = c_i^{\text{init}} + \sum_{t=1}^{T} \frac{m_{it}}{(1+r)^t} + \frac{d_i}{(1+r)^T} \qquad (1)$$

where c_i^{init} is initial expenditure, m_{it} is recurring expenditure, d_i is net retirement expenditure, and r is a consistently chosen discount rate. This is an accounting convention, not an estimate of a security return. Existing sunk expenditure is excluded from the incremental comparison. Shared expenditure must be assigned once, through a common enabling action or a documented allocation.

A present-value budget does not ensure annual cash availability. Where annual funding matters, constraints $c_{it}x_i \leq B_t$ are required in addition to, or instead of, the aggregate limit. Similarly, the aggregate staff capacity below is appropriate only where effort is sufficiently interchangeable within the period. Separate constraints are needed for scarce skills, recurring operations, or time windows. The simple form makes the essential distinction visible without asserting that all organizations can pool effort freely.

Supplier promises are evidence inputs, not demonstrated outcomes. NIST supply-chain guidance addresses risks associated with acquired products and services [18]. The proposed application is to retain support commitments, access arrangements, acceptance tests, and renewal costs with each candidate. If a critical contribution cannot be substantiated, the decision should include an evidence-gathering action or a conservative scenario, rather than silently assigning the promised benefit.

D. Values and Uncertainty Assumptions

Let Ω be a finite set of assessment scenarios. Each describes a plausible set of conditions relevant to the same protection service. A scenario may emphasize credential misuse, unauthorized physical entry, or service interruption. It can also represent pessimistic assumptions about how well a proposed control works. It is not automatically an equally probable event in a stochastic model.

For each scenario, decision participants assess an incremental value a_i^{ω} for action i. Values are cardinal stakeholder assessments on a common interval scale. Equal point differences must represent equal assessed differences in the decision objective, including across scenarios. An ordinal label such as low, medium, or high cannot simply be substituted for these coefficients. Where a common cardinal interpretation cannot be agreed, the proposed objective is unsuitable; a non-compensatory comparison or a partially ordered shortlist should be used instead.

One possible elicitation decomposes the service into distinct outcome attributes, agrees value functions and trade-offs, and assesses action contributions relative to the baseline. The same outcome anchors and trade-offs must be retained across scenarios: independently normalizing each scenario would change the maximin comparison. Zero denotes no incremental contribution relative to the fixed baseline; point totals need not be bounded by 100. Independent technical review checks whether the same protected consequence has been credited

more than once. Complementary actions can receive an additional joint contribution b_{ij}^{ω}, representing only the increment beyond their individual contributions. Its mechanism and evidence must be recorded. Substitutable or overlapping measures should first be reformulated into non-overlapping alternatives or constrained as alternatives. The model below does not claim to represent arbitrary higher-order interactions.

These values are not probabilities, annual losses, percentages of incidents prevented, or monetary returns. Their credibility depends on the elicitation and evidence process. Scenario analysis exposes disagreement and uncertainty; it does not remove them. Adding an omitted adverse scenario can change a recommendation that previously looked robust.

E. Finite-Scenario Portfolio Formulation

Let N contain the candidate actions. Binary variable x_i equals one when action i is selected. Let M contain mandatory actions, P contain directed prerequisite pairs (i,j) meaning that i requires j, Q contain mutually exclusive pairs, and E contain pairs with an assessed complementary contribution. Let B be the incremental funding limit and H the available staff capacity. Cost c_i and effort h_i use consistent units within their respective constraints.

For $(i,j) \in E$, binary variable z_{ij} indicates joint selection. Scenario value is

$$V_{\omega}(x,z) = \sum_{i\in N} a_i^{\omega}\, x_i + \sum_{(i,j)\in E} b_{ij}^{\omega}\, z_{ij} \qquad (2)$$

The following mixed-integer linear program maximizes the lowest assessed value over the stated scenarios:

$$\max_{x,z,\eta}\ \eta \qquad (3)$$

$$\text{s.t. } \eta \le V_{\omega}(x,z),\ \ \omega \in \Omega \qquad (4)$$

$$\sum_{i\in N} c_i\, x_i \le B,\ \ \sum_{i\in N} h_i\, x_i \le H \qquad (5)$$

$$x_i = 1,\ \ i \in M \qquad (6)$$

$$x_i \le x_j,\ \ (i,j) \in P \qquad (7)$$

$$x_i + x_j \le 1,\ \ (i,j) \in Q \qquad (8)$$

$$z_{ij} \le x_i,\ \ z_{ij} \le x_j,$$

$$z_{ij} \ge x_i + x_j - 1,\ \ (i,j) \in E \qquad (9)$$

$$x_i, z_{ij} \in \{0,1\},\ \ \eta \in \mathbb{R} \qquad (10)$$

Constraints (9) enforce $z_{ij} = x_i x_j$, so a complement cannot receive credit when one member is absent. Prerequisites and complementarity have different meanings: a prerequisite establishes admissibility; a complement changes the assessed value of an admissible combination. A mandatory action may also be a prerequisite for another action.

The primary objective treats all included scenarios as conditions against which the portfolio is assessed, without probability weights. Among primary optima, the decision rule minimizes total cost and then staff effort. Any remaining tie is reported; a fixed identifier order is used only for reproducibility of the illustration. Unspent budget is permitted. If a funded portfolio cannot meet the mandatory requirements and their prerequisites, the model reports infeasibility. It does not relax those requirements automatically or describe an infeasible portfolio as an acceptable minimum.

Observation 1. For a fixed action set, scenario set, coefficients, and mandatory conditions, the optimal minimum value cannot decrease when either resource limit is increased while the other is unchanged, provided the original problem is feasible.

Proof. Increasing a resource limit retains every previously feasible portfolio. Maximization over this enlarged feasible set cannot yield a smaller objective value. This establishes weak monotonicity, not a strictly positive marginal value of additional resources. ▫

The distinction matters for budgeting. A funding increase may leave the chosen portfolio unchanged because staff effort, a prerequisite, or an indivisible action remains binding. Ratios of individual value to cost cannot generally express such interactions. Moreover, a favorable model value is a statement about the specified decision representation. It is not proof that the portfolio will deliver the intended protection.

F. Decision Record and Review

The portfolio calculation is followed by technical and managerial review. The record should identify the service boundary, baseline, mandatory conditions, horizon, cost sources, capacity assumptions, scenario definitions, coefficient rationale, interactions, feasible alternatives, and sensitivity findings. It should also explain why each deferred action was deferred and name the person authorized to accept the remaining exposure. An attractive optimum is not self-authorizing.

Implementation acceptance should check the function that motivated the expenditure. Purchasing backup power is different from demonstrating that the relevant service continues during a controlled outage. NIST's measurement guidance supports linking measures to information needs [19]. Here, proposed follow-up evidence includes accepted functional tests, actual internal effort, recurring costs, and deviations from the assumed dependencies. No universal targets are specified.

G. Design of the Computational Illustration

The computational illustration uses six invented actions, three invented scenarios, arbitrary planning units, and stipulated value points. All 64 binary selections are enumerated. Feasible portfolios are compared under the same assumptions using a maximin rule, a scenario-specific objective, and a precisely specified ratio heuristic. This verifies the calculations and illustrates decision consequences. It does not validate real-world control effectiveness, stakeholder preferences, costs, or avoided losses.

The illustration concerns a fictional facility with a fixed existing baseline. Six actions are listed in Table II. All numbers are invented for analytical demonstration. Costs are arbitrary lifecycle planning units, not currency or market prices; staff effort is expressed in arbitrary capacity units. Values use one stipulated interval scale across three assessment scenarios: ω_1 emphasizes unauthorized physical entry, ω_2 emphasizes privileged credential misuse, and ω_3 emphasizes interruption of the protection service. The numbers are not calibrated estimates of the effectiveness of the named measures.

TABLE II. CONSTRUCTED INPUTS FOR THE SIX-ACTION ILLUSTRATION

Action	Illustrative scope	c_i	h_i	a_i^1	a_i^2	a_i^3
A	Baseline service and safety review	3	2	2	2	2
B	Privileged access hardening	4	3	2	6	2
C	Door and controller maintenance	4	3	6	1	3
D	Power and local fallback	5	3	3	2	7
E	Joint monitoring integration	4	4	4	4	2
F	Cross-domain response exercise	2	2	2	2	3

Action A is mandatory in this constructed decision. This is a stipulated planning requirement, not a legal claim. Action E requires B: $x_E \leq x_B$. Actions C and F have a joint incremental contribution $(b^1_{CF}, b^2_{CF}, b^3_{CF}) = (3,1,2)$. All other pair contributions are zero, and Q is empty. Every action has a non-negative incremental contribution in every scenario. The pair bonus is separate from the standalone entries. Its exact value is an assumption to be challenged, not a literature-derived estimate.

The main resource limits are $B = 14$ and $H = 10$. The same constraints and scenario values apply to all comparisons. The robust rule maximizes $\min_\omega V_\omega$. The scenario-specific rule maximizes V_1; it does not maximize an expected value. Both use the cost and effort tie-breaks described above.

The ratio heuristic first selects A. It then considers each remaining action once, in descending order of $\min_\omega a_i^\omega / c_i$, breaking equal ratios by identifier. An action is added only if it fits both remaining resources and its prerequisite is already selected. Skipped actions are not revisited. The resulting consideration order is F, B, E, D, C. This benchmark is deliberately simple and fully specified. It does not stand for every possible greedy or multicriteria method.

Exhaustive enumeration inspects all $2^6 = 64$ selections and retains 15 feasible portfolios under the main resource limits. It provides an exact optimum for this small instance without a heuristic optimality claim. For larger instances, formulation (3)–(10) can be passed to a mixed-integer solver, with feasibility status, optimality gap, and time limit reported. No scalability experiment is presented here.

Because the instance is deterministic and every admissible selection is evaluated exactly, the reported values are computed quantities rather than sample estimates. Measures of sample size, dispersion, and statistical significance therefore do not apply. The enumeration was implemented in Python 3 using only the standard library.

IV. RESULTS

A. Portfolio Comparison

Table III reports the computed portfolios. The robust rule selects ABCF, with cost 13, effort 10, and scenario values (15,12,12). For example, its value in scenario 2 is $2 + 6 + 1 + 2 + 1 = 12$, including the C–F contribution. The ratio heuristic selects ABDF. The scenario-specific optimum is ACDF, which has the largest value in scenario 1 but a lower value under credential misuse.

TABLE III. EXACT RESULTS FOR THE CONSTRUCTED INSTANCE WITH B = 14 AND H = 10

Rule	Portfolio	Cost	Effort	V_1	V_2	V_3	Min.
Robust portfolio	ABCF	13	10	15	12	12	12
Ratio heuristic	ABDF	14	10	9	12	14	9
Scenario 1 optimum	ACDF	14	10	16	8	17	8

Relative to ACDF, the robust portfolio ABCF has one point less in scenario 1 and a minimum value of 12 rather than 8. ACDF retains a higher value in scenario 3, as does ABDF. The ratio heuristic selects D before reaching C and does not obtain the joint C–F increment. ABCF uses 13 of the 14 available cost units and all 10 staff units.

B. Sensitivity to Funding, Capacity, and Dependencies

Table IV reports resource and structural sensitivity checks. Structural changes are compared with the original assumptions at the same resource limits. With $H = 10$, increasing the funding limit from 13 to 19 does not improve the optimal minimum value. Increasing funding alone is therefore unproductive within this model and range. At $B = 14$, reducing staff capacity to 8 changes the optimum to ABC and the minimum to 7. Capacity of 6 yields AF with a minimum of 4. These changes reflect constructed constraints, not measured staffing requirements of real controls.

TABLE IV. RESOURCE AND STRUCTURAL SENSITIVITY RESULTS FOR THE HYPOTHETICAL INSTANCE

Changed condition	Portfolio	Min.	Cost	Effort
$B = 9, H = 10$	ABF	6	9	7
$B = 11, H = 10$	ABC	7	11	8
$B = 13, H = 10$	ABCF	12	13	10
$B = 19, H = 10$	ABCF	12	13	10
$B = 14, H = 8$	ABC	7	11	8
$B = 14, H = 6$	AF	4	5	4
No C–F bonus; $B = 14, H = 10$	ABCF	10	13	10
Prerequisite omitted; $B = 9, H = 10$	AEF	7	9	8

Removing the C–F bonus leaves ABCF selected at $B = 14$, $H = 10$, but changes its scenario values to (12,11,10). The final row deliberately deletes $x_E \leq x_B$ from the model at $B = 9$, $H = 10$. The resulting recommendation is AEF, with minimum value 7, whereas the correctly constrained optimum is ABF with minimum 6. Finally, since A alone costs 3 and requires effort 2, a budget below 3 or a capacity below 2 makes the instance infeasible.

V. DISCUSSION

This study asked how physical and cybersecurity investments can be prioritized under budget and staffing constraints when controls are interdependent and outcome estimates are uncertain. The results for the constructed instance are consistent with the working proposition stated in Section I: the individual ratio ordering did not identify the feasible portfolio with the best worst-case assessed value. Because the instance is hypothetical, this demonstrates a mathematical possibility rather than an empirical regularity.

A. Interpretation of the Illustrative Results

The robust portfolio sacrifices one point in scenario 1 relative to ACDF while increasing its minimum value from 8 to 12 within this stipulated scale. It does not dominate ACDF, which also has a higher value in scenario 3; ABDF likewise retains a higher scenario 3 value than ABCF. The comparison exposes a preference for protection across all declared conditions, rather than establishing that one portfolio is universally better. The meaning of the robust choice rests on acceptance of the objective and scenario set.

The ratio heuristic does not recognize the joint C–F increment when sorting standalone actions, and it commits resources to D before reaching C. This instance therefore supplies a reproducible counterexample to the claim that an individual ratio ordering necessarily solves the constrained portfolio problem. It establishes no general performance bound for ratio methods. It also shows why portfolio selection can leave budget unused: the preferred ABCF uses 13 of 14 cost units and exhausts the staff envelope.

The continued selection of ABCF without the C–F bonus is informative only about that one perturbation. It does not establish robustness to all elicitation errors, combined changes, or omitted scenarios. The recommendation AEF obtained when the prerequisite is deleted is inadmissible under the stated implementation relationship because B is missing;

its apparently improved result comes from a misspecified feasible set. The lesson is structural: numerical optimization cannot compensate for an omitted prerequisite. Similarly, when the mandatory action cannot be funded, the appropriate output is an escalation decision concerning resources, scope, or the authorized baseline. Quietly discarding A would solve a different problem.

B. What the Procedure Adds to a Security Budget

The procedure makes a distinction that departmental budgets often obscure. A risk priority identifies a concern requiring attention; an investment portfolio identifies an implementable combination of responses. Ranking alone does not account for shared prerequisites, indivisible actions, or the staff effort required to make a control operational. The constructed example demonstrates these differences without claiming that its coefficients transfer to another setting.

Mandatory conditions also change the meaning of optimization. Their treatment as constraints prevents a large discretionary benefit from offsetting failure to meet an authorized minimum. This is consistent with retaining a separate safety decision rather than attempting to value every consequence in one compensatory score. The substantive basis and authority for each mandatory condition still require review. Overusing the mandatory label can pre-empt deliberation just as readily as omitting an essential condition can distort it.

The maximin objective addresses uncertainty through a preference for the least favorable assessed value within the declared scenarios. It can be conservative. An implausible or poorly defined extreme scenario may dominate the result, while an omitted credible scenario may make the result overconfident. Decision-makers should inspect which scenarios bind and compare alternative objective conventions. If probabilities become defensible, an expected-value formulation may be appropriate, potentially alongside minimum-service constraints. The present analysis does not estimate such probabilities.

C. Implementation and Accountability

A practical pilot should begin with a bounded service and a manageable set of actions. Physical security and facilities define operating requirements and fallback conditions; technical specialists examine digital and infrastructure dependencies; finance reviews comparable cost treatment; and business leadership approves trade-offs. Supplier participation can clarify delivery assumptions, but vendors should not determine the value of their own proposals without independent scrutiny.

The decision record should name an accountable owner for each deferred exposure. Deferral may reflect an unavailable specialist, insufficient evidence, a prerequisite scheduled first, or a deliberate acceptance decision. These reasons have different review triggers. Treating all of them as a single unfunded backlog conceals information that matters at the next budget round.

The review cycle should compare predicted inputs with delivered capability. Relevant evidence includes realized expenditure, internal effort, acceptance failures, continuing support costs, and the response observed during authorized exercises. A completed purchase order is evidence of expenditure; a successful functional check provides different evidence about service capability. Neither, by itself, establishes a causal reduction in incidents.

An evidence-gathering action can be an investment candidate. Examples include checking the operating state of a controller, confirming the feasibility of backup restoration, or obtaining a maintenance estimate. Its value should not be invented to avoid loss. A more developed model could represent a staged decision in which learning changes later choices. The single-stage model here simply identifies uncertainty and the circumstances under which a new decision is needed.

D. Limitations and Validation Agenda

The literature review is targeted and not exhaustive. Differences in access to complete articles limit the detail of some comparisons. Evidence concerning SMEs, industrial systems, and supply chains informs the decision structure but cannot establish how frequently its problems occur in the facilities considered here. The framework is intended to be proportionate to resources; it has not been validated for a particular size class, industry, or jurisdiction.

The model assumes a fixed baseline, a single planning stage, known resource requirements within each calculation, a common cardinal value scale, and only pairwise complementary contributions. Actual resource requirements may be uncertain, higher-order interactions may matter, and an attacker may adapt to investment. Shared controls can also produce overlap that an additive representation fails to capture. Treating their contributions as independent would overstate value. Excluding or bundling overlapping alternatives is a modelling choice requiring substantive review, not a universal solution.

The staff constraint is aggregated, and the illustration has no scheduling or annual cash-flow structure. Equation (1) does not imply that discounted affordability ensures liquidity. Data collection for an operational application must therefore distinguish planning value, cash availability, implementation effort, and continuing maintenance. No return on investment, annualized loss reduction, or optimal share of revenue is estimated in this study.

Validation should first test whether multidisciplinary participants can identify mandatory conditions and dependencies consistently. A second stage should examine elicitation reliability, including whether participants interpret value differences comparably and whether double counting is detected. Authorized pilot applications could then compare recommendations with feasible alternatives and record actual delivery effort and functional acceptance. A longitudinal study would be needed to investigate realized security outcomes while accounting for changes in exposure, systems, staffing, and reporting. The exact solution of a hypothetical instance is not a substitute for those stages.

VI. CONCLUSION

Prioritizing physical and cybersecurity under limited budgets is a constrained decision about protection capabilities. Individual risk rankings or value-to-cost ratios cannot generally determine a feasible portfolio when actions require prerequisites, compete for staff, or produce complementary value. A credible decision process must also distinguish mandatory service conditions from discretionary improvements and retain uncertainty about the evidence supporting each action.

This paper provides an auditable procedure and a finite-scenario binary formulation connecting those requirements. The hypothetical example demonstrates the logic: the selected robust portfolio differs from both a specified ratio heuristic and a scenario-specific optimum; extra funding has no value

over part of the examined range; and an omitted prerequisite produces an inadmissible recommendation. These conclusions follow from the declared mathematical instance and do not establish operational effectiveness.

The study has limitations. The literature review is targeted rather than exhaustive, and the illustration is hypothetical and uncalibrated. The model assumes a fixed baseline, a single planning stage, an aggregated staff constraint, a common cardinal value scale, and only pairwise complementary contributions.

The practical contribution is a decision record linking service objectives, candidate actions, resource limits, uncertainty, selected combinations, and residual-risk ownership. Its usefulness in real organizations remains to be evaluated through expert review, elicitation studies, and observed implementation. The model supports accountable judgment; it does not replace the evidence needed to justify a security investment.

DATA AVAILABILITY

All numerical inputs are hypothetical and are reported in Tables II–IV. No organizational, client, or participant data were collected. The enumeration script, written in Python 3 using only the standard library, is available from the author upon reasonable request.

REFERENCES

[1] L. A. Gordon and M. P. Loeb, "The economics of information security investment," *ACM Trans. Inf. Syst. Secur.*, vol. 5, no. 4, pp. 438–457, 2002, doi: 10.1145/581271.581274.

[2] L. A. Gordon, M. P. Loeb, and L. Zhou, "Integrating cost–benefit analysis into the NIST Cybersecurity Framework via the Gordon–Loeb model," *J. Cybersecurity*, vol. 6, no. 1, Art. no. tyaa005, 2020, doi: 10.1093/cybsec/tyaa005.

[3] M. Heidt, J. P. Gerlach, and P. Buxmann, "Investigating the security divide between SME and large companies: How SME characteristics influence organizational IT security investments," *Inf. Syst. Front.*, vol. 21, no. 6, pp. 1285–1305, 2019, doi: 10.1007/s10796-019-09959-1.

[4] A. Chidukwani, S. Zander, and P. Koutsakis, "A survey on the cyber security of small-to-medium businesses: Challenges, research focus and recommendations," *IEEE Access*, vol. 10, pp. 85701–85719, 2022, doi: 10.1109/ACCESS.2022.3197899.

[5] C. Rombaldo Junior, I. Becker, and S. D. Johnson, "Unaware, unfunded, untrained, and unsupported: A systematic review of cybersecurity in small and medium enterprises," *J. Cybersecurity*, vol. 12, no. 1, Art. no. tyag030, 2026, doi: 10.1093/cybsec/tyag030.

[6] H. I. Kure, S. Islam, and M. A. Razzaque, "An integrated cyber security risk management approach for a cyber-physical system," *Appl. Sci.*, vol. 8, no. 6, Art. no. 898, 2018, doi: 10.3390/app8060898.

[7] T. Sawik and B. Sawik, "A rough cut cybersecurity investment using portfolio of security controls with maximum cybersecurity value," *Int. J. Prod. Res.*, vol. 60, no. 21, pp. 6556–6572, 2022, doi: 10.1080/00207543.2021.1994166.

[8] A. Sukumar, H. A. Mahdiraji, and V. Jafari-Sadeghi, "Cyber risk assessment in small and medium-sized enterprises: A multilevel decision-making approach for small e-tailors," *Risk Anal.*, vol. 43, no. 10, pp. 2082–2098, 2023, doi: 10.1111/risa.14092.

[9] Y. Cherdantseva *et al.*, "A review of cyber security risk assessment methods for SCADA systems," *Comput. Secur.*, vol. 56, pp. 1–27, 2016, doi: 10.1016/j.cose.2015.09.009.

[10] P. Kobis, "Human factor aspects in information security management in the traditional IT and cloud computing models," *Oper. Res. Decis.*, vol. 31, no. 1, pp. 61–76, 2021, doi: 10.37190/ord210104.

[11] National Institute of Standards and Technology, "The NIST Cybersecurity Framework (CSF) 2.0," NIST, Gaithersburg, MD, USA, CSWP 29, 2024, doi: 10.6028/NIST.CSWP.29.

[12] S. Quinn *et al.*, "Integrating cybersecurity and enterprise risk management (ERM)," NIST, Gaithersburg, MD, USA, IR 8286 Rev. 1, 2025, doi: 10.6028/NIST.IR.8286r1.

[13] Joint Task Force Transformation Initiative, "Guide for conducting risk assessments," NIST, Gaithersburg, MD, USA, SP 800-30 Rev. 1, 2012, doi: 10.6028/NIST.SP.800-30r1.

[14] D. Bertsimas and M. Sim, "The price of robustness," *Oper. Res.*, vol. 52, no. 1, pp. 35–53, 2004, doi: 10.1287/opre.1030.0065.

[15] H. Snyder, "Literature review as a research methodology: An overview and guidelines," *J. Bus. Res.*, vol. 104, pp. 333–339, 2019, doi: 10.1016/j.jbusres.2019.07.039.

[16] S. Kriaa, L. Piètre-Cambacédès, M. Bouissou, and Y. Halgand, "A survey of approaches combining safety and security for industrial control systems," *Reliab. Eng. Syst. Saf.*, vol. 139, pp. 156–178, 2015, doi: 10.1016/j.ress.2015.02.008.

[17] K. Stouffer *et al.*, "Guide to operational technology (OT) security," NIST, Gaithersburg, MD, USA, SP 800-82 Rev. 3, 2023, doi: 10.6028/NIST.SP.800-82r3.

[18] J. Boyens, A. Smith, N. Bartol, K. Winkler, A. Holbrook, and M. Fallon, "Cybersecurity supply chain risk management practices for systems and organizations," NIST, Gaithersburg, MD, USA, SP 800-161 Rev. 1 (updated Nov. 2024), 2022, doi: 10.6028/NIST.SP.800-161r1-upd1.

[19] K. Schroeder, H. Trinh, and V. Pillitteri, "Measurement guide for information security: Volume 1—Identifying and selecting measures," NIST, Gaithersburg, MD, USA, SP 800-55 Vol. 1, 2024, doi: 10.6028/NIST.SP.800-55v1.