Business Risk Analytics for Insider Threat Prevention in Data-Driven Organizations: An Entropy-Based Decision Framework

Shubing Hou[1, *], Zhaoqi Wang[1], Qingyin Wei[1]

[1] School of Management, Capital Normal University, Beijing, China 100048

* Email: 13733176685@163.com (Corresponding Author)

Received 30 March 2026 · Revised 14 May 2026 · Accepted 28 June 2026 · Published 30 September 2026

Abstract

Insider-related incidents can undermine data-intensive organizations even when access originates from legitimate accounts and authorized workflows. Security teams must therefore decide where to intervene while distinguishing exposure from proof of misconduct and balancing protection, business continuity, employee privacy, and limited resources. This article develops a business risk analytics framework that combines workflow-level indicators with Shannon entropy weighting and a separate economic decision stage. Five indicators represent sensitive-data exposure, excessive access, outbound-control gaps, operational dependence, and detection or response gaps. Fixed reference scales preserve the interpretation of individual indicators, whereas entropy weights describe their dispersion across the workflows under comparison; they do not estimate incident probabilities or causal importance. A completely synthetic four-workflow illustration yields entropy weights of 0.182, 0.402, 0.080, 0.148, and 0.188, respectively, and ranks financial reporting, software delivery, customer analytics, and procurement from highest to lowest composite priority. Equal weights reverse the middle pair, and a plausible change in one synthetic input also reverses that pair. A separate hypothetical budget example shows how independently assumed avoided losses and implementation costs can inform control selection. These calculations establish arithmetic transparency but do not demonstrate detection accuracy, causal effectiveness, or financial returns in an operating organization. The framework adds explicit measurement rules, non-compensatory safeguards, uncertainty analysis, privacy protections, and a prospective validation protocol to entropy-based prioritization. It offers a transparent basis for managerial discussion while requiring empirical calibration and independent governance before operational use.

Keywords: Insider threat; Business risk analytics; Entropy weighting; Multicriteria decision analysis; Security governance; Privacy; Organizational resilience

Business Risk Analytics for Insider Threat Prevention in Data-Driven Organizations: An Entropy-Based Decision Framework

1. Introduction

Organizations now depend on integrated business applications, cloud repositories, analytics platforms, and collaborative workflows to create and use valuable data. These arrangements allow employees, contractors, service providers, and automated identities to move information across systems in support of legitimate work. The same access can enable unauthorized disclosure, manipulation, disruption, or accidental exposure. A compromised account may also appear to be acting from within the organization, although the ultimate actor is external. Surveys of insider threats show that threat sources, motives, available evidence, and potential countermeasures are heterogeneous (Homoliak et al., 2020; Liu et al., 2018). For managers, this diversity makes a single technical alert or employee label an insufficient basis for allocating preventive resources.

Insider-related risk has at least three analytically distinct forms: intentional misuse of entrusted access, unintentional action that exposes an asset, and misuse of a legitimate identity by an outside actor. These categories can overlap in observed logs and may not be distinguishable at the beginning of an investigation. An unusually large transfer, for example, may represent a legitimate analytics job, a configuration error, an account compromise, or a deliberate attempt to remove data. Concluding intent from the transfer alone would be unsound. Human-factor research emphasizes that time pressure, task design, and the surrounding organizational environment affect the conditions under which inadvertent incidents arise (Khan et al., 2022; Pathirana et al., 2026).

Organizations also encounter a business allocation problem. Access reviews, secure transfer controls, monitoring improvements, backup arrangements, and staff support all compete for budgets and operational attention. A technically detectable anomaly is not automatically a high-consequence event, and a process with a serious exposure may have few alerts because telemetry is absent. Existing cybersecurity risk work highlights the need to connect threat, vulnerability, and consequence with managerial preferences and available interventions (Ganin et al., 2020; Shameli-Sendi et al., 2016). Yet the practical translation from heterogeneous workflow evidence to a defensible sequence of decisions remains difficult when reliable incident frequencies and loss distributions are sparse.

This article asks how an organization can transparently prioritize preventive review across business workflows when indicators differ in units, incident labels are incomplete, and decision makers must account for employee privacy. Its aim is not to build a classifier that identifies malicious individuals. The unit of analysis is a business workflow within a defined assessment period, such as financial reporting or software delivery. The proposed index summarizes exposure and control gaps for those workflows; any investigation of a specific event requires a separate, evidence-based process. Restricting the unit in this way enables management to discuss risk drivers and control ownership without treating a numerical score as a statement about a person's character.

The proposed framework makes four contributions. First, it translates five measurable dimensions of exposure and resilience into a common, predeclared scale, with explicit treatment of missing data and access limitations. Second, it uses entropy to summarize variation across the selected workflows while separating statistical discrimination from managerial significance. Third, it introduces non-compensatory safeguards and a second-stage control-allocation model in which costs and avoided losses are independently estimated. Fourth, it specifies sensitivity analyses and an empirical evaluation design, including negative-case review, privacy safeguards, and monitoring for changes in the data. The numerical example is synthetic and intended to make the procedure reproducible; it is not evidence that the approach improves security or produces a return on investment.

The remainder of the article reviews insider-threat and organizational security research, sets out the decision framework and its assumptions, works through a synthetic illustration, and considers control selection, implementation, and validation. The discussion identifies the conditions under which an entropy-weighted index may support informed choice and the circumstances in which additional evidence or alternative methods are necessary.

2. Literature Review and Research Positioning

2.1 Insider threats and organizational evidence

Insider-threat research covers a range of actors and activity patterns rather than a single class of malicious employee. The survey by Al-Mhiqani et al. (2020) organizes detection approaches, data sources, and open problems, while Gheyas and Abdallah (2016) review prediction and detection studies and draw attention to variation in modeling approaches. Inayat et al. (2024) examine prevention and mitigation strategies, and Alsowail and Al-Shehari (2021) propose a layered prevention framework. Taken together, these studies suggest that prevention requires coordinated controls across data, access, workflows, and response rather than the assumption that one signal can establish both intent and impact.

Technical research increasingly models relations among identities, devices, assets, and events. Gong et al. (2024) review graph-based detection; Cai et al. (2024) develop a model using adaptive neighbors; and Li et al. (2024) investigate few-shot detection using graph meta-learning. Roy and Chen (2024) consider cyber-human relationships, while Fei et al. (2025) study the analysis of log text. These approaches can reveal structure that simple counts miss. Their use nevertheless depends on data quality, access permissions, and the validity of the labels and environments used to test them. In this article, such technical systems are potential upstream sources of aggregate indicators, not proof that the proposed workflow index has detection capability.

The operational setting matters. Erola et al. (2022) report lessons from deploying an insider-threat detection tool in multinational organizations, showing why integration and practical performance deserve attention alongside algorithmic results. A survey of explainable artificial intelligence for insider detection considers the need to make detection outputs interpretable (Saeed Alketbi & Mehmood, 2025). Explanation, however, must identify the inputs, comparisons, and limits of a decision; a plausible narrative attached to an alert cannot establish that the underlying inference is correct. At the workflow level, the present framework therefore retains each component and its provenance so that a reviewer can distinguish a sensitive asset from a poorly measured control gap.

2.2 Human factors, compliance, and organizational context

Security behavior is shaped by both individual beliefs and organizational arrangements. Bulgurcu et al. (2010) examine employee awareness and rationality-based beliefs about information security policy compliance. Herath and Rao (2009) consider penalties, social pressures, and perceived effectiveness, while D'Arcy et al. (2009) examine awareness of security countermeasures and information-systems misuse. Chen et al. (2012) analyze different approaches to organizational compliance. These studies support the view that rules and controls need a credible organizational context; they do not imply that noncompliance can be inferred reliably from a person's demographic or psychological characteristics.

Additional behavioral work highlights potential unintended effects of security demands. D'Arcy et al. (2014) investigate employee responses to stressful security requirements, Cram et al. (2021) study information security fatigue, and Trang and Nastjuk (2021) examine stress and policy design in an experimental setting. Dang-Pham et al. (2025) discuss shadow security practices, including workarounds associated with employees' attempts to accomplish tasks. The relevance to prevention is practical: an apparent gap may reflect an unusable process, conflicting responsibilities, or insufficient support. A control response that ignores such

context could shift work outside approved systems or impose burdens without reducing the underlying exposure.

The literature also considers leadership, motivation, and the design of awareness measures. Hu et al. (2012) examine management participation and organizational culture, while Menard et al. (2017) compare motivational perspectives on security behavior. Parsons et al. (2017) report further validation of a human-aspects questionnaire, and Sohrabi Safa et al. (2016) propose an organizational compliance model. Crossler et al. (2013) and Delso-Vicente et al. (2025) review directions and factors in behavioral security research. These sources justify examining the work environment and control design. They do not justify inserting unvalidated personality scores or private communications into an individual risk index.

2.3 Multicriteria risk decisions and information entropy

Multicriteria decision analysis (MCDA) structures choices involving several measures that cannot be compared directly. Ganin et al. (2020) frame cybersecurity management as a multicriteria decision problem, and Maček et al. (2020) review applications of MCDA to information security risk assessment. Some security controls interact, so evaluating them in isolation may misrepresent their collective effect (Lo & Chen, 2012). Ou Yang et al. (2013) also examine interrelationships among information-security risk controls in a decision framework. These contributions motivate an explicit separation between prioritizing where to investigate and selecting which controls to implement.

Shannon (1948) introduced entropy as a measure of uncertainty in a probability distribution. In entropy weighting, normalized shares of observed criterion values are used to describe how evenly a criterion is distributed across the alternatives. Zhu et al. (2020) discuss entropy weighting in decision making, while Roszkowska (2026) examines the importance of normalization for multicriteria rankings. The mathematical similarity does not mean that a workflow's normalized share is an empirical probability of an insider incident. Here, entropy measures dispersion in the decision matrix, and the resulting weight measures relative discrimination within that matrix. The approach makes weights reproducible conditional on the data, but it cannot supply social values, causal effects, or a cost of loss by itself.

Prioritization research also emphasizes the need to support the people who act on analytical outputs. Jalalvand et al. (2025) survey alert prioritization criteria and methods, and Baruwal Chhetri et al. (2024) explore human–AI teaming for alert fatigue. Chatziamanetoglou and Rantos (2025) develop weighted criteria for assessing cyber threat intelligence quality. These studies motivate a decision record that is intelligible to security analysts and business owners, with an explicit distinction between the value of the evidence and the seriousness of the exposed asset. The proposed approach builds on these themes at the workflow level and adds a concrete protocol for protecting against misleading weights and poorly supported economic conclusions.

3. Research Design and Decision Framework

3.1 Scope, decision unit, and process

This is a conceptual and methodological study supported by a transparent synthetic numerical illustration. It is not a systematic literature review, an empirical field study, or an operational evaluation. The proposed decision process begins with the definition of a set of business workflows and a common assessment window. Each workflow should have an accountable owner, a documented service boundary, an inventory of data and identities, and an agreed business purpose. A workflow may cross several applications, but the analyst must define how shared assets and events are allocated so that a single gap is not counted repeatedly without explanation.

The framework proceeds through six stages: establish permissible data collection and mandatory protections; define indicators and reference scales; assemble and quality-check a workflow-by-indicator matrix; calculate

entropy weights and a priority index; compare alternative weights and inputs; and select controls using separate evidence on cost, benefit, and feasibility. A governance group comprising the business owner, security personnel, privacy representatives, and other authorized stakeholders reviews the result. It records overrides and their reasons. The process is periodic rather than continuously self-adjusting, so that an unstable or incomplete data feed does not silently change organizational priorities.

The analytical priority is a screening aid. It is not an estimate of the likelihood, severity, or expected financial loss from an actual incident, and it is not a basis for disciplinary action. If the organization needs a monetary risk measure, it must estimate event frequency and conditional consequence with appropriate uncertainty, a different and more demanding task. Likewise, a high index may arise from poor controls or high business dependence even where no anomalous behavior has been observed. Treating that distinction as a design principle allows managers to improve a workflow without stigmatizing the people who use it.

3.2 Indicator definitions and fixed reference scales

Table 1 proposes five indicators, each expressed on a 0–1 scale where higher values indicate greater exposure or weaker preparedness. Organizations should document a measurement dictionary and use the same definitions across workflows and assessment periods. In the example, x_1, x_2, x_3, and x_5 are ratios whose numerator is a subset of a clearly specified denominator; x_4 is a scaled business-impact estimate. The data source, time window, owner, calculation rules, and quality checks must be recorded for every cell. A comparable scale is useful, but it does not make the indicators equivalent measures of risk.

Table 1. Proposed Workflow-Level Indicators

Indicator	Dimension	Definition on a 0–1 scale
x_1	Sensitive-data exposure	Fraction of in-scope records classified as confidential or regulated.
x_2	Excess privilege	Fraction of active in-scope identities with verified excess access.
x_3	Transfer-control gap	Fraction of in-scope outbound channels lacking an effective required control.
x_4	Continuity consequence	Estimated loss from a specified 24-hour interruption divided by a fixed reference amount, capped at 1.
x_5	Detection and response gap	Fraction of essential telemetry or response checks missing, stale, or not triaged within the agreed window.

All definitions and values require an organization-specific measurement dictionary.

For x_1, the asset inventory should define the population of records and a policy-based rule for classifying confidential or regulated records. A fraction does not fully represent the value of a small but exceptionally sensitive collection; such assets may require a separate mandatory review. For x_2, an excess privilege is an entitlement that remains active beyond the approved role and has been verified against an up-to-date entitlement catalog. A reviewer should distinguish a true exception from approved emergency access. Shared and service accounts belong in the defined identity population if the organization can assign ownership and assess them consistently.

For x_3, the denominator is the number of inventoried channels or interfaces through which the workflow can transfer covered data, not the number of all network packets. A channel is counted as a gap when the specified transfer control is absent or ineffective, according to a documented test. For x_4, the business owner estimates the consequence of a specified interruption, such as a 24-hour loss of the workflow, and divides it by a predeclared reference amount, clipping values above one. The consequence estimate may include direct service costs and recovery work, but its scope must be common to all workflows. It is a scenario impact, not expected annual loss.

For x_5, an organization should define the essential telemetry and response checks for a workflow before calculating the fraction that is absent, stale, or cannot be triaged within the agreed window. This measure exposes a visibility gap; a missing log is not an observed malicious event. A low number of alerts must not be interpreted as strong security when monitoring coverage is low. If a denominator is zero, unknown, or unreliable, the cell is marked missing and referred to its owner rather than replaced with zero. When a workflow genuinely has no relevant outbound channel, the indicator is structurally inapplicable and a separate assessment design is required.

A workflow inventory should also capture dependencies that fall outside its immediate owner's control. Customer analytics, for example, may rely on a shared identity provider, an external processing service, and a central data warehouse. The assessment team should state whether a common control is represented once at enterprise level or allocated to each workflow according to a consistent rule. It should avoid treating repeated entries for the same weakness as independent incidents. Conversely, an asset used by several workflows may require a joint response because fixing it for one process does not necessarily protect the others. These boundary choices should be reviewed whenever a service, supplier, or business function changes.

The data dictionary should specify the observation window, the latest acceptable timestamp, and tests for duplicate accounts, assets, and events. Source systems may undercount dormant accounts or miss activity outside the approved perimeter. A comparison is meaningful only when the workflows are measured against compatible coverage and the same reference anchors. If evidence is delayed, the matrix should include a freshness field; data quality can be displayed separately rather than disguised as a favorable indicator. Before scoring, the governance group may set minimum completeness requirements and route workflows with critical unknowns to manual review.

3.3 Entropy weights and the composite index

Let $i = 1, \ldots, m$ index workflows and $j = 1, \ldots, n$ index indicators, with $m \geq 2$ and $x_{ij} \in [0, 1]$. All indicators are oriented so that larger values describe greater exposure or control weakness. For each indicator with a positive column sum, define its share, normalized entropy, and degree of dispersion as follows:

$$p_{ij} = x_{ij} / \Sigma_i x_{ij} \qquad H_j = -[\Sigma_i p_{ij} \ln(p_{ij})] / \ln(m) \qquad d_j = 1 - H_j \qquad (1)$$

Here ln is the natural logarithm and the convention $0 \ln 0 = 0$ is used. The quantity p_{ij} is the share of the observed indicator total assigned to workflow i; it is not an incident probability. If every entry in a column is zero, define $H_j = 1$ and $d_j = 0$. A column with the same positive value in every workflow likewise has $H_j = 1$ and zero dispersion. These rules avoid division by zero and correctly show that the indicator cannot distinguish the compared workflows, even though a common level of exposure can be high.

If at least one indicator has nonzero dispersion, the entropy weight and composite index are:

$$w_j = d_j / \Sigma_j d_j \qquad R_i = \Sigma_j w_j x_{ij} \qquad (2)$$

The score is a weighted average on the original anchored 0–1 scales. With fixed nonnegative weights that sum to one, an increase in any indicator cannot reduce the same workflow's score. This local monotonicity does not necessarily hold when the weights are recalculated after a change in the full matrix: altering one value can change every weight and consequently the scores of other workflows. The weights are relative to the selected comparison set, and adding or removing a workflow can change them. Organizations should therefore preserve the set definition and weight version in every assessment record.

When all indicators are constant, the sum of d_j is zero and entropy weights are undefined. The proposed fallback is to use an explicitly documented equal-weight index for descriptive reporting; all workflows then have identical component values and identical scores. The correct managerial response is to inspect absolute exposures and mandatory controls, not to infer that common exposure is harmless. For near-constant data,

tiny differences or measurement noise can dominate the weight calculation. A minimum quality and dispersion threshold, fixed before the review, should trigger an equal-weight or expert-weight comparison and an explanation rather than automatic reliance on an unstable result.

A hybrid alternative incorporates a separately documented set of policy weights v_j, where $v_j \geq 0$ and the weights sum to one. The analyst selects a mixing parameter λ between zero and one and computes $w_j(\lambda) = \lambda w_j + (1 - \lambda)v_j$. The entropy and policy components must be reported separately because the choice of λ expresses governance preferences. To assess robustness, the organization can compare pure entropy, equal weights, and several plausible hybrid values. No particular λ is claimed to be universally optimal. Ranges can be used when an indicator is known only approximately, but the assumptions and dependency among ranges must be recorded.

3.4 Non-compensatory safeguards and uncertainty

A weighted average permits compensation: a very high value in one dimension can be offset by low values in others. That behavior is inappropriate for mandatory protections. Before ranking, the organization should apply preapproved gates for conditions such as unowned privileged accounts, missing controls over a legally sensitive dataset, or an untested critical recovery capability. A breached gate requires a designated response regardless of the composite score. The specific thresholds and obligations depend on the organization and governing law and must be confirmed locally; they are not derived from entropy.

Uncertainty arises from measurement error, a changing workflow population, uncertain impact estimates, and disagreement over control effectiveness. The assessment should therefore preserve source evidence, assess each cell's completeness and confidence, and repeat the calculation over defensible alternative values. Important scenarios include a lower and higher business-impact estimate, an updated entitlement catalog, alternative normalization anchors, and exclusion of a workflow. If alternatives change priority, the organization should report that instability and investigate the influential inputs. A narrow numerical difference should not be presented as an objectively meaningful ordering without a measurement error analysis.

Correlated indicators introduce another problem. For example, a sensitive-data fraction and a business-impact measure may both reflect the same underlying asset value. Counting them separately can make the asset appear more important than intended. Analysts should review definitions for conceptual overlap and examine correlations when the sample is adequate, but correlation estimates from four workflows are not reliable evidence of redundancy. If indicators are combined or removed, the new version should be approved and the full sensitivity comparison reported. A structured index makes these choices visible; it does not remove the need for judgment.

4. Synthetic Numerical Illustration

4.1 Input matrix and computation

The example is entirely invented for demonstration. It contains no organizational records, observed incidents, employee data, externally calibrated parameters, or measured control effects. Four workflows are labeled A, financial reporting; B, customer analytics; C, software delivery; and D, procurement. Five synthetic indicator values are shown in Table 2. The values should be read as assumed results of the measurement rules in Table 1, not as estimates for any real organization. Their purpose is to show how the calculations and sensitivity tests work.

Table 2. Synthetic Indicator Matrix

Workflow	x_1	x_2	x_3	x_4	x_5
A Financial reporting	0.95	0.85	0.65	0.80	0.60

Workflow	x_1	x_2	x_3	x_4	x_5
B Customer analytics	0.90	0.60	0.80	0.55	0.70
C Software delivery	0.75	0.90	0.55	0.85	0.45
D Procurement	0.55	0.40	0.70	0.60	0.80

All inputs are hypothetical and are not derived from an organization.

For the first indicator, the sum of values is 3.15. The four shares are therefore approximately 0.3016, 0.2857, 0.2381, and 0.1746. Substitution into Eq. (1) gives $H_1 = 0.9853$ and $d_1 = 0.0147$. Repeating the calculation for the other columns yields the entropies, dispersion values, and normalized weights in Table 3. For example, the second indicator has the highest weight because its assumed values differ most strongly in proportional terms within this particular set. It does not follow that excessive privileges are always the most consequential insider-threat factor.

Table 3. Entropy, Dispersion, and Weights for the Synthetic Matrix

Indicator	Entropy H_j	Dispersion d_j	Weight w_j
x_1	0.9853	0.0147	0.1817
x_2	0.9674	0.0326	0.4021
x_3	0.9935	0.0065	0.0797
x_4	0.9880	0.0120	0.1483
x_5	0.9847	0.0153	0.1882

Calculations use unrounded intermediate values; displayed results are rounded to four decimals.

Using unrounded values in Eq. (2), the workflow scores are A = 0.7978, B = 0.6818, C = 0.7528, and D = 0.5561. The entropy-weighted priority order is A, C, B, D. A is highly exposed across multiple indicators; C's high privilege and continuity values have substantial influence because the second indicator receives the largest entropy weight. D receives a lower composite score despite its high detection and response gap. This does not mean that D's monitoring gap is acceptable, and it does not relieve its owner from a mandatory control requirement.

4.2 Comparison and sensitivity

Equal weights produce scores of 0.7700, 0.7100, 0.7000, and 0.6100 for A through D, respectively. The resulting order is A, B, C, D: customer analytics and software delivery swap positions. This comparison demonstrates that a statement such as “C is more risky than B” depends on how the five dimensions are valued. The differences between the two methods are not a test of which ranking is correct. Both are conditional summaries of the assumed matrix, and neither has been compared with actual future losses in this study.

A second check changes only D's privilege-gap value from 0.40 to 0.80, a hypothetical correction to the underlying evidence. Recalculating all entropy weights gives approximately 0.2613, 0.1401, 0.1146, 0.2133, and 0.2707. The corresponding scores become A = 0.7749, B = 0.7177, C = 0.6882, and D = 0.6806. The order becomes A, B, C, D. Although B and C's own inputs did not change, their relative rank reverses because the distribution and weights changed. This is a consequence of the method and a reason to monitor data revisions.

A third calculation estimates weights using A, B, and C only, then applies those weights to all four original rows for comparison. The resulting weights are approximately 0.0767, 0.2289, 0.1845, 0.2621, and 0.2479. The full-set scores under these weights are A = 0.7457, B = 0.6716, C = 0.6993, and D = 0.6184; the order

remains A, C, B, D. Although the ordering is unchanged, the weights and scores differ materially. In practice, D's excluded row would be an out-of-sample case for weight calculation, and the comparison set should be justified rather than chosen to obtain a preferred result.

These tests expose two distinct sources of uncertainty: disagreement about how indicators should be weighted and sensitivity to the observed comparison set. A manager should consider B and C a potentially unstable middle group rather than claim that one is definitively safer. The analysis also illustrates why unrounded intermediate values must be retained for calculations. Table 3 reports rounded numbers for readability; rerunning the formulas with the values in Table 2 reproduces the results within rounding tolerance. No predictive accuracy, reduction in incidents, or real-world effectiveness is established by this arithmetic exercise.

5. Translating Priorities into Control Decisions

5.1 The economic decision is a separate model

The score identifies where further review may be warranted; it does not identify the most valuable intervention. A control's value depends on the scenario it addresses, its effectiveness, its cost, and interactions with existing defenses. Gordon and Loeb (2002) analyze information-security investment under uncertainty, illustrating the need to relate spending to potential losses and the nature of vulnerability. In practice, evidence about control effects may be sparse and managerial estimates may be subjective. The assessment should distinguish an empirically estimated benefit from an expert judgment and make the assumed time horizon explicit.

Let z_k equal one if discretionary control k is selected and zero otherwise. If annual avoided losses b_k and annualized control costs c_k can be estimated on a consistent basis, and the benefits are non-overlapping, the organization can maximize $\Sigma(b_k - c_k)z_k$ subject to $\Sigma c_k z_k \leq B$. Here B is the available annualized discretionary budget. The model is a simple binary allocation problem; it assumes that the listed costs include relevant implementation and operating expenses and that the consequences of the selected controls are additive. It does not include the workflow index as a substitute for b_k. Mandatory controls are funded and handled separately before this discretionary budget is considered.

Consider three hypothetical options: privilege reconciliation, controlled data transfer, and task-oriented coaching. Suppose their respective annual costs in thousands of dollars are 30, 40, and 12, while their separately assumed annual avoided losses are 45, 60, and 20. With an annualized discretionary budget of 80, the first two controls cost 70 and yield an assumed net benefit of 35. The second and third cost 52 and yield 28; the first and third cost 42 and yield 23. Selecting all three would cost 82 and exceed the budget. Under these assumptions, the first two are the best feasible combination among the listed choices. The amounts are invented for the illustration and are neither empirical estimates nor a demonstrated return on investment.

The calculation changes if benefits overlap. For example, a secure transfer control and privilege reduction may prevent the same disclosure, so adding their gross benefits would count that avoided loss twice. An expanded model can specify the value of each control portfolio or include interaction terms, but each added quantity must be justified. If benefits are uncertain, analysts can compare portfolios over low, central, and high scenarios, or estimate distributions when sufficient data exist. A constrained decision may also require minimum service availability, implementation capacity, privacy safeguards, and contractual obligations. A portfolio that appears attractive on expected net benefit alone may fail those requirements.

5.2 From findings to action

For a high x_2 value, the first response is to verify the identity inventory and the basis for each entitlement. The outcome may be removal of obsolete access, a time-limited elevation mechanism, or an improved approval process. In a software-delivery workflow, emergency access and automated deployment identities

may be essential to continuity; a blanket removal could cause an outage. A business owner and an access administrator should therefore agree on exceptions and recovery procedures. The reduction in excess entitlement after review is a process measure, not proof that a malicious incident was prevented.

For a high x_3 value, a control review can inspect the destinations, permitted file types, classification labels, and operational need for each outbound channel. A business unit may need a secure partner exchange rather than a prohibition on all transfers. For a high x_4 value, protection can include recovery testing, segregation of duties, and continuity arrangements. For a high x_5 value, the organization should establish whether logs are missing, whether events are usable, and whether an analyst can respond in time. Each intervention needs an owner, a completion date, and a verification method. Retesting the same indicators may show that a control gap has closed, but additional evidence is required to establish a causal reduction in losses.

An escalation path should document how an authorized analyst distinguishes workflow-level control review from investigation of a particular event. If a plausible incident emerges, the analyst preserves relevant evidence and follows an approved incident-response procedure with proportionate access and independent review. If the event is explained by a legitimate task, the record should preserve that disposition and allow the workflow or detection rule to be corrected. Employees need a way to report mistakes and contest factual errors without assuming that all unusual activity is culpable. Such procedures make the business decision actionable while limiting the risk of converting an uncertain screening signal into an unsupported allegation.

6. Implementation, Evaluation, and Governance

6.1 Implementation and data quality

Implementation begins with a pilot in a small number of comparable workflows selected before the matrix is calculated. The organization should record why the workflows were included, which business processes were excluded, and whether a workflow has unusual monitoring coverage. Each source system should produce a versioned extract with the observation period, identifier, owner, and completeness checks. The analyst should reconcile counts with the asset inventory and investigate changes caused by system migrations, reorganizations, or reclassification. Fenz and Ekelhart (2011) distinguish verification, validation, and evaluation in information-security risk management; each is needed here because correct arithmetic does not establish that the measures represent the intended constructs.

Verification checks whether the implementation reproduces the stated formulas and boundary rules. It should include zero-valued columns, constant columns, identical rows, the $m = 2$ boundary, incomplete data, and rounding. Validation asks whether the indicators and reference values reasonably represent the exposures they claim to measure. A process owner might review a stratified sample of excessive-access findings or verify whether supposedly unprotected transfer paths actually carry sensitive data. Evaluation asks whether using the framework improves an operational decision under specified conditions. These are separate questions and require different evidence; a successful software test alone answers only the first.

Data quality checks should cover denominator consistency, duplicates, the time since last refresh, and the proportion of systems actually observed. Where values are missing, the preferred response is to recover the evidence or hold the affected workflow for review. If imputation is necessary for research, it should be declared, based only on information available at the assessment time, and accompanied by a sensitivity analysis. Replacing missing monitoring data with a zero exposure score is particularly hazardous because it rewards invisibility. A confidence flag or interval can accompany the index, but any threshold for withholding a score must be agreed before the assessment.

6.2 A prospective validation design

A prospective evaluation should define a use case and comparator before collecting outcomes. One possible design compares an entropy-assisted workflow review with the organization's existing risk register or an

equal-weight index. The primary outcome could be the number of independently verified, materially relevant control gaps corrected per unit of review effort, observed over a fixed follow-up period. Secondary outcomes might include time to correction, business-process disruption caused by controls, privacy complaints, and the cost of review. The outcome definition must state who verifies a gap and what evidence qualifies as correction; simply counting alerts or changed scores is inadequate.

For an intervention study, eligible workflows could be allocated to review sequences in a randomized or otherwise defensible design, subject to operational and ethical constraints. Where random allocation is impractical, a matched comparison or interrupted time-series analysis may be considered, with explicit attention to changes in staffing, technology, and external threat conditions. Baseline, follow-up, selection criteria, and planned analyses should be specified in advance. A statistical association between receiving a review and a lower incident count does not by itself establish that the framework caused the reduction, particularly when severe workflows receive more attention and more incidents are detected as a result.

The assessment should also be compared with the practical decisions that managers would otherwise make. A prospective log can capture the original priority, the recommended control, the reason for accepting or rejecting it, the resources committed, and the eventual verification result. Independent reviewers can assess a sample of these records without seeing which method produced each recommendation. Alongside the primary outcome, researchers should report how much evidence gathering each approach requires and whether the comparison caused delays in normal business operations. This design can reveal whether an apparently more detailed index actually improves resource allocation or merely increases documentation.

Any study of incident detection requires a defensible event definition and a way to examine both reviewed and unreviewed cases. Labels created only from high-score workflows create verification bias: unreviewed workflows are treated as safe because no one looked. An authorized team can instead review a documented sample across score levels, including low-score or non-alert cases, subject to privacy and labor requirements. Assessment data must precede the outcomes they are intended to predict. When data are divided for model selection and testing, time-respecting splits and separation of related identities or workflows help reduce leakage from future activity or repeated records.

If a separate classifier is introduced for event triage, evaluation should report precision, recall, the number of alerts per analyst, and performance at the actual operating threshold, together with uncertainty and the underlying event prevalence. The base-rate problem can make even an apparently accurate system burdensome (Axelsson, 2000); precision–recall analysis is particularly informative when positive cases are rare (Saito & Rehmsmeier, 2015). As an illustrative calculation, suppose 100 of 100,000 events are truly positive, sensitivity is 90%, and the false-positive rate is 2%. The system detects 90 positive events but flags 1,998 negative events, so precision is 90/2,088, approximately 4.3%. These figures are hypothetical and do not describe the proposed index or an actual deployment.

The study should also assess stability across sites, business functions, employment arrangements, and time periods when such comparisons are lawful and sufficiently supported. Small groups or poorly verified labels can make subgroup estimates misleading or identifiable. Differences in alert rates may arise from role requirements, monitoring coverage, or scheduling rather than wrongdoing. Investigators should examine these mechanisms and compare error rates only when adequate ground truth exists. A fairness review should include the consequences of an investigation, not merely a parity statistic for the initial score. Findings that are too uncertain to support comparison should be reported as such.

6.3 Privacy, proportionality, and accountability

The framework is designed to operate on aggregate workflow measures wherever possible. It should not require private message content, inferred mental state, demographic attributes, or individual character assessments to calculate its five indicators. Privacy choices affect trust and perceived trade-offs; Kehr et al.

(2015) examine institutional trust and general privacy concerns in the privacy calculus. Before deployment, organizations should specify the purpose of collection, the legal and employment basis for processing, permitted access, retention periods, and routes for correction. They should assess alternatives that achieve the same control objective with less intrusive information.

Access to raw identity-level evidence should be restricted to an authorized role and separated from the business-level dashboard. The dashboard should show component values, uncertainty, data freshness, and mandatory-control status without exposing unnecessary personal details. Investigative access should require a documented trigger and be auditable. Where people can be affected by an outcome, the organization should establish appropriate notice, review, appeal, and correction procedures in accordance with applicable requirements. Automated scores must not be used as sole evidence for discipline. The framework itself supplies no legal authorization and cannot resolve jurisdiction-specific labor or privacy obligations.

The operating model should name a control owner, a data steward, an independent reviewer, and a person accountable for approving changes to the indicator dictionary and weight procedure. Monthly or quarterly review cycles can be chosen according to business conditions, but the interval is a governance choice rather than a validated prescription. A review should record actions taken, exceptions, unresolved data problems, and the rationale for any management override. If weights or indicators are changed, the old results and calculation version should remain available for audit. The organization should investigate abrupt shifts before using them to reallocate resources.

7. Discussion

7.1 Theoretical and managerial implications

The framework connects sociotechnical insider-risk research with a practical multicriteria decision sequence. Its central proposition is that observable control conditions, exposure of valuable data, and operational dependence can be considered together without treating them as direct observations of intent. The technical literature supplies examples of what may be detectable, while organizational research explains why workflow design and incentives matter. The analytical contribution is the explicit linkage among fixed indicator definitions, entropy-based discrimination, non-compensatory gates, and a separately justified economic choice.

Entropy weighting can be attractive because another analyst can reproduce the weights from the same matrix. Its apparent objectivity is nevertheless conditional on choices about the population of workflows, the observation period, the indicator scale, and the measurement system. Fixed anchors avoid the direct rescaling of every value when a new alternative becomes the minimum or maximum, but they do not eliminate the dependence of entropy weights on the comparison set. A high weight says that a criterion distinguishes the observed workflows more strongly under the specified calculation. It does not show that an intervention on that criterion produces the greatest reduction in harm.

The synthetic example has a useful managerial interpretation. Financial reporting ranks first under both the entropy and equal-weight calculations, but the ordering of the middle workflows is not robust. A sensible committee could therefore retain A as a candidate for early review while examining B and C together, subject to any mandatory control failures. The evidence on D's high monitoring gap warrants its own investigation even though D ranks last in the composite. A documented override is not a failure of the method when it follows from a predeclared rule or newly verified information; it can reveal where the model needs a better indicator or a different decision boundary.

7.2 Limitations and future research

The illustration contains only four invented workflows. Small matrices make entropy weights sensitive to individual values, and the demonstration provides no basis for estimating performance in a population. The

indicators may not capture collusion, changing attacker tactics, vendor dependencies, or a low-frequency event with catastrophic consequences. A single time window also misses sequences and evolving exposure. Deployments should be compared with existing assessment processes before adoption, and indicators should be adapted to the setting while preserving the conceptual distinction between exposure, evidence quality, and actual loss.

The meaning of an improvement also depends on the evaluation period. A privilege review can reduce recorded excessive entitlements shortly after implementation, while its effect on future incidents may emerge slowly and may be difficult to observe because serious incidents are rare. A longer follow-up can strengthen evidence but introduces changes in staffing, applications, and attacker behavior. Researchers should distinguish a successfully implemented control from a verified reduction in exposure, and both from a demonstrated change in realized losses. Reporting these outcomes separately allows a decision maker to understand what has actually been observed and what remains an assumption.

The proposed additive index cannot express every interdependency among controls and may obscure high-severity conditions unless the gate layer is enforced. Expert-defined impact anchors and policy weights introduce judgment that may vary among stakeholders. Similarly, the budget example relies on known, additive benefits and costs; real interventions can have delayed effects, joint benefits, and uncertain opportunity costs. Future work should study alternative weighting procedures, portfolio interactions, interval-based decision methods, and approaches that explicitly represent rare severe outcomes. Comparative studies should report their assumptions and baselines rather than select a method solely because it gives a preferred ranking.

Research access to organizational evidence may be constrained by privacy, security, and confidentiality. Synthetic data can support code verification and exploratory design, but cannot establish that a model generalizes to real behavior. A meaningful evaluation would require lawful access, independent assessment of outcomes, representative negative cases, transparent handling of uncertainty, and the ability to report unfavorable results. Research on organizational controls should also examine whether additional monitoring changes employee behavior, workload, or trust. These outcomes are integral to the business decision rather than secondary effects to be ignored.

8. Conclusion

This article developed an entropy-based decision framework for organizing workflow-level insider-risk information in data-driven organizations. It defined five interpretable indicators, a transparent weighting and scoring procedure, mandatory safeguards, sensitivity tests, and a separate economic allocation stage. In the hypothetical matrix, the rank of two workflows changed under equal weights and after one input correction, illustrating that prioritization depends on modeling choices and evidence quality. A separate hypothetical budget calculation showed how control selection can be framed when costs and non-overlapping benefits are independently assumed. The analysis does not establish detection accuracy or demonstrate organizational benefit. Its practical value must be assessed through prospective, privacy-conscious research and careful implementation in the intended setting.

Acknowledgement

We acknowledge the researchers whose published work informed this study. Any remaining errors are our responsibility.

Author Contributions

Conceptualization, S.B.H.; Methodology, S.B.H.; Formal Analysis, S.B.H.; Investigation, S.B.H. and Z.Q.W. and Q.Y.W.; Data Curation, S.B.H. and Z.Q.W. and Q.Y.W.; Software, S.B.H. and Z.Q.W. and Q.Y.W.; Validation, S.B.H. and Z.Q.W. and Q.Y.W.; Resources, Z.Q.W. and Q.Y.W.; Visualization, S.B.H. and Z.Q.W. and Q.Y.W.; Writing – Original Draft, S.B.H.; Writing – Review & Editing, S.B.H. and Z.Q.W. and Q.Y.W.; Supervision, S.B.H.; Project Administration, S.B.H. We have read and approved the final manuscript and take responsibility for the work.

Funding

This research received no external funding.

Conflicts of Interest

We declare no conflicts of interest.

Ethical Approval and Consent

The illustrative calculations use synthetic values and involve no human participants or personal records. Any future study involving organizational or personal data should obtain the approvals and permissions applicable to that study.

Data Availability Statement

All synthetic inputs and the formulas required to reproduce the numerical illustration are reported in this article. No real organizational dataset was used.

Declaration of Generative AI Use

Generative AI assisted us with drafting and editorial preparation. We take full responsibility for the accuracy of the content and references and for the final manuscript.

References

Al-Mhiqani, M. N., Ahmad, R., Zainal Abidin, Z., Yassin, W., Hassan, A., Abdulkareem, K. H., Ali, N. S., & Yunos, Z. (2020). A review of insider threat detection: Classification, machine learning techniques, datasets, open challenges, and recommendations. Applied Sciences, 10(15), 5208. https://doi.org/10.3390/app10155208

Alsowail, R. A., & Al-Shehari, T. (2021). A multi-tiered framework for insider threat prevention. Electronics, 10(9), 1005. https://doi.org/10.3390/electronics10091005

Axelsson, S. (2000). The base-rate fallacy and the difficulty of intrusion detection. ACM Transactions on Information and System Security, 3(3), 186–205. https://doi.org/10.1145/357830.357849

Baruwal Chhetri, M., Tariq, S., Singh, R., Jalalvand, F., Paris, C., & Nepal, S. (2024). Towards human-AI teaming to mitigate alert fatigue in security operations centres. ACM Transactions on Internet Technology, 24(3), 1–22. https://doi.org/10.1145/3670009

Bulgurcu, B., Cavusoglu, H., & Benbasat, I. (2010). Information security policy compliance: An empirical study of rationality-based beliefs and information security awareness. MIS Quarterly, 34(3), 523–548. https://doi.org/10.2307/25750690

Cai, X., Wang, Y., Xu, S., Li, H., Zhang, Y., Liu, Z., & Yuan, X. (2024). LAN: Learning adaptive neighbors for real-time insider threat detection. IEEE Transactions on Information Forensics and Security, 19, 10157–10172. https://doi.org/10.1109/tifs.2024.3488527

Chatziamanetoglou, D., & Rantos, K. (2025). Weighted quality criteria for cyber threat intelligence: Assessment and prioritisation in the MISP data model. International Journal of Information Security, 24(4), 160. https://doi.org/10.1007/s10207-025-01080-6

Chen, Y., Ramamurthy, K., & Wen, K.-W. (2012). Organizations' information security policy compliance: Stick or carrot approach? Journal of Management Information Systems, 29(3), 157–188. https://doi.org/10.2753/mis0742-1222290305

Cram, W. A., Proudfoot, J. G., & D'Arcy, J. (2021). When enough is enough: Investigating the antecedents and consequences of information security fatigue. Information Systems Journal, 31(4), 521–549. https://doi.org/10.1111/isj.12319

Crossler, R. E., Johnston, A. C., Lowry, P. B., Hu, Q., Warkentin, M., & Baskerville, R. (2013). Future directions for behavioral information security research. Computers & Security, 32, 90–101. https://doi.org/10.1016/j.cose.2012.09.010

Dang-Pham, D., Thompson, N., Ahmad, A., & Maynard, S. (2025). Shadow information security practices in organizations: The role of information security transparency, overload, and psychological empowerment. Computers & Security, 156, 104538. https://doi.org/10.1016/j.cose.2025.104538

D'Arcy, J., Herath, T., & Shoss, M. K. (2014). Understanding employee responses to stressful information security requirements: A coping perspective. Journal of Management Information Systems, 31(2), 285–318. https://doi.org/10.2753/mis0742-1222310210

D'Arcy, J., Hovav, A., & Galletta, D. (2009). User awareness of security countermeasures and its impact on information systems misuse: A deterrence approach. Information Systems Research, 20(1), 79–98. https://doi.org/10.1287/isre.1070.0160

Delso-Vicente, A.-T., Diaz-Marcos, L., Aguado-Tevar, O., & de Blanes-Sebastián, M. G. (2025). Factors influencing employee compliance with information security policies: A systematic literature review of behavioral and technological aspects in cybersecurity. Future Business Journal, 11(1), 28. https://doi.org/10.1186/s43093-025-00452-7

Erola, A., Agrafiotis, I., Goldsmith, M., & Creese, S. (2022). Insider-threat detection: Lessons from deploying the CITD tool in three multinational organisations. Journal of Information Security and Applications, 67, 103167. https://doi.org/10.1016/j.jisa.2022.103167

Fei, K., Zhou, J., Zhou, Y., Gu, X., Fan, H., Li, B., Wang, W., & Chen, Y. (2025). LaAeb: A comprehensive log-text analysis based approach for insider threat detection. Computers & Security, 148, 104126. https://doi.org/10.1016/j.cose.2024.104126

Fenz, S., & Ekelhart, A. (2011). Verification, validation, and evaluation in information security risk management. IEEE Security & Privacy Magazine, 9(2), 58–65. https://doi.org/10.1109/msp.2010.117

Ganin, A. A., Quach, P., Panwar, M., Collier, Z. A., Keisler, J. M., Marchese, D., & Linkov, I. (2020). Multicriteria decision framework for cybersecurity risk assessment and management. Risk Analysis, 40(1), 183–199. https://doi.org/10.1111/risa.12891

Gheyas, I. A., & Abdallah, A. E. (2016). Detection and prediction of insider threats to cyber security: A systematic literature review and meta-analysis. Big Data Analytics, 1(1), 6. https://doi.org/10.1186/s41044-016-0006-0

Gong, Y., Cui, S., Liu, S., Jiang, B., Dong, C., & Lu, Z. (2024). Graph-based insider threat detection: A survey. Computer Networks, 254, 110757. https://doi.org/10.1016/j.comnet.2024.110757

Gordon, L. A., & Loeb, M. P. (2002). The economics of information security investment. ACM Transactions on Information and System Security, 5(4), 438–457. https://doi.org/10.1145/581271.581274

Herath, T., & Rao, H. R. (2009). Encouraging information security behaviors in organizations: Role of penalties, pressures and perceived effectiveness. Decision Support Systems, 47(2), 154–165. https://doi.org/10.1016/j.dss.2009.02.005

Homoliak, I., Toffalini, F., Guarnizo, J., Elovici, Y., & Ochoa, M. (2020). Insight into insiders and IT: A survey of insider threat taxonomies, analysis, modeling, and countermeasures. ACM Computing Surveys, 52(2), 30. https://doi.org/10.1145/3303771

Hu, Q., Dinev, T., Hart, P., & Cooke, D. (2012). Managing employee compliance with information security policies: The critical role of top management and organizational culture. Decision Sciences, 43(4), 615–660. https://doi.org/10.1111/j.1540-5915.2012.00361.x

Inayat, U., Farzan, M., Mahmood, S., Zia, M. F., Hussain, S., & Pallonetto, F. (2024). Insider threat mitigation: Systematic literature review. Ain Shams Engineering Journal, 15(12), 103068. https://doi.org/10.1016/j.asej.2024.103068

Jalalvand, F., Baruwal Chhetri, M., Nepal, S., & Paris, C. (2025). Alert prioritisation in security operations centres: A systematic survey on criteria and methods. ACM Computing Surveys, 57(2), 42. https://doi.org/10.1145/3695462

Kehr, F., Kowatsch, T., Wentzel, D., & Fleisch, E. (2015). Blissfully ignorant: The effects of general privacy concerns, general institutional trust, and affect in the privacy calculus. Information Systems Journal, 25(6), 607–635. https://doi.org/10.1111/isj.12062

Khan, N., Houghton, R. J., & Sharples, S. (2022). Understanding factors that influence unintentional insider threat: A framework to counteract unintentional risks. Cognition, Technology & Work, 24(3), 393–421. https://doi.org/10.1007/s10111-021-00690-z

Li, X., Li, L., Li, X., Cai, B., Jia, J., Gao, Y., & Yu, S. (2024). GMFITD: Graph meta-learning for effective few-shot insider threat detection. IEEE Transactions on Information Forensics and Security, 19, 7161–7175. https://doi.org/10.1109/tifs.2024.3430106

Liu, L., De Vel, O., Han, Q.-L., Zhang, J., & Xiang, Y. (2018). Detecting and preventing cyber insider threats: A survey. IEEE Communications Surveys & Tutorials, 20(2), 1397–1417. https://doi.org/10.1109/comst.2018.2800740

Lo, C.-C., & Chen, W.-J. (2012). A hybrid information security risk assessment procedure considering interdependences between controls. Expert Systems with Applications, 39(1), 247–257. https://doi.org/10.1016/j.eswa.2011.07.015

Maček, D., Magdalenić, I., & Ređep, N. (2020). A systematic literature review on the application of multicriteria decision making methods for information security risk assessment. International Journal of Safety and Security Engineering, 10(2), 161–174. https://doi.org/10.18280/ijsse.100202

Menard, P., Bott, G. J., & Crossler, R. E. (2017). User motivations in protecting information security: Protection motivation theory versus self-determination theory. Journal of Management Information Systems, 34(4), 1203–1230. https://doi.org/10.1080/07421222.2017.1394083

Ou Yang, Y.-P., Shieh, H.-M., & Tzeng, G.-H. (2013). A VIKOR technique based on DEMATEL and ANP for information security risk control assessment. Information Sciences, 232, 482–500. https://doi.org/10.1016/j.ins.2011.09.012

Parsons, K., Calic, D., Pattinson, M., Butavicius, M., McCormac, A., & Zwaans, T. (2017). The human aspects of information security questionnaire (HAIS-Q): Two further validation studies. Computers & Security, 66, 40–51. https://doi.org/10.1016/j.cose.2017.01.004

Pathirana, N., Roberts, R., Kalutarage, H., & McDermott, C. D. (2026). Integrating human factors into insider threat detection – a systematic review. ACM Computing Surveys, 58(10), 260. https://doi.org/10.1145/3798089

Roszkowska, E. (2026). Entropy and normalization in MCDA: A data-driven perspective on ranking stability. Entropy, 28(1), 114. https://doi.org/10.3390/e28010114

Roy, K. C., & Chen, G. (2024). GraphCH: A deep framework for assessing cyber-human aspects in insider threat detection. IEEE Transactions on Dependable and Secure Computing, 21(5), 4495–4509. https://doi.org/10.1109/tdsc.2024.3353929

Saeed Alketbi, K., & Mehmood, A. (2025). A comprehensive survey of explainable artificial intelligence techniques for malicious insider threat detection. IEEE Access, 13, 121772–121798. https://doi.org/10.1109/access.2025.3587114

Saito, T., & Rehmsmeier, M. (2015). The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets. PLOS ONE, 10(3), e0118432. https://doi.org/10.1371/journal.pone.0118432

Shameli-Sendi, A., Aghababaei-Barzegar, R., & Cheriet, M. (2016). Taxonomy of information security risk assessment (ISRA). Computers & Security, 57, 14–30. https://doi.org/10.1016/j.cose.2015.11.001

Shannon, C. E. (1948). A mathematical theory of communication. Bell System Technical Journal, 27(3), 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

Sohrabi Safa, N., Von Solms, R., & Furnell, S. (2016). Information security policy compliance model in organizations. Computers & Security, 56, 70–82. https://doi.org/10.1016/j.cose.2015.10.006

Trang, S., & Nastjuk, I. (2021). Examining the role of stress and information security policy design in information security compliance behaviour: An experimental study of in-task behaviour. Computers & Security, 104, 102222. https://doi.org/10.1016/j.cose.2021.102222

Zhu, Y., Tian, D., & Yan, F. (2020). Effectiveness of entropy weight method in decision-making. Mathematical Problems in Engineering, 2020, 3564835. https://doi.org/10.1155/2020/3564835