

Publicly Disclosed Corporate Security Governance in Critical-Service Firms: Accountability Structures and Disclosure-Specificity Signals

Karim Hardy

hardykarim@gmail.com

École Nationale Supérieure des Mines de Paris

Research Article

Keywords: corporate security governance, public disclosure, enterprise risk management, third-party governance, critical-service firms, content analysis

Posted Date: October 6th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10494511/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License. [Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Public reporting is often used to infer corporate security governance, although documentary specificity cannot establish practice. We analysed 31 FY2025 public documents from 24 listed EU/EEA critical-service firms using a ten-domain, evidence-linked 0–3 coding scheme. Enterprise-risk-management integration was the most consistently disclosed governed signal (23/24 at code 3), whereas third-party governance was least often disclosed at the named-mechanism threshold (16/24 at ≥ 2 ; five non-disclosures). Three transparent, non-exclusive bundles recurred, but no stable mutually exclusive typology was retained. Evidence was predominantly cyber/information-security mediated: excluding explicitly cyber-only records left 91/240 evaluable cells. Results describe public-disclosure specificity and auditability only; they do not measure implementation, organizational maturity, compliance, resilience, or operational effectiveness.

1. Introduction

Firms that provide energy, transport, digital connectivity, banking, and financial-market infrastructure occupy a consequential position between private organization and public continuity. European policy reflects that position by assigning critical and essential entities explicit responsibilities for risk management, resilience, incident reporting, and supply-chain security (European Parliament and Council 2022a, 2022b). Yet external observers rarely have direct access to how those responsibilities are implemented. What they can inspect are annual reports, governance statements, sustainability reports, security policies, and other public documents. These materials make some accountability structures visible, but visibility is not equivalent to control quality or operational performance.

The distinction matters because corporate disclosure serves several purposes at once. It reduces selected information asymmetries, responds to institutional expectations, and communicates what management considers material or legitimate to report (Healy and Palepu 2001; Verrecchia 2001). Risk-reporting research also shows that annual reports vary in amount, specificity, and narrative content (Linsley and Shrives 2006). A named board committee, executive owner, supplier requirement, or assurance mechanism is therefore a traceable public signal. It is not direct evidence that the mechanism operates consistently, reaches all relevant units, or reduces loss. Treating the two as interchangeable would turn a documentary study into an unsupported assessment of organizational practice.

Relevant scholarship is substantial but fragmented across adjacent objects. Enterprise risk management (ERM) studies examine implementation, risk cultures, and integration into strategy and control (Beasley, Clune, and Hermanson 2005; Mikes 2009). Information-security governance research emphasizes accountability, oversight, and alignment (Williams 2001; Posthumus and von Solms 2004). Supply-chain research addresses third-party dependencies and cybersecurity exposure (Boyson 2014). Corporate risk-disclosure studies focus on narrative reporting and determinants of disclosure. These strands illuminate elements of corporate security governance, but they do not by themselves provide an evidence-linked, cross-domain account of what critical-service firms make publicly visible.

The residual gap is methodological as well as substantive. Existing disclosure research demonstrates how corporate narratives can be coded systematically (Abraham and Cox 2007; Deumes 2008), while content-analysis methodology requires explicit units, categories, decision rules, and evidence traceability (Krippendorff 2018). What remains underdeveloped is a bounded design that: (1) distinguishes generic commitments, named mechanisms, and governed signals; (2) retains the firm-by-domain profile rather than publishing a league table; (3) tests transparent overlapping bundles without forcing firms into types; and (4) reports how source scope, uncertainty, and cyber-dominant evidence change the conclusions.

This gap is consequential for both research and review. A reader can usually verify whether a report names a responsible body, connects a risk to an enterprise process, or describes an assurance route. The same reader cannot infer from that statement whether meetings occurred as described, controls covered the relevant assets, or corrective actions changed outcomes. Documentary evidence is therefore strongest when the construct and the inference are coextensive: the observable claim concerns what the document makes auditable, while implementation remains outside the claim (Bowen 2009; Adcock and Collier 2001).

The study asks: How do publicly listed European critical-service firms disclose the allocation, integration, and assurance of corporate security governance, and what recurring configurations of disclosed accountability and disclosure-specificity signals emerge across sectors? Four subquestions identify which governance mechanisms are publicly named; assess the specificity and auditability of signals across ten domains; examine recurring cross-domain configurations and descriptive subgroup variation; and test sensitivity to document scope, ordinal thresholds, and uncertain cases. Here, “disclosure-specificity signals” refers only to graded public-disclosure specificity. It does not denote internal capability or actual operational maturity.

To answer those questions, the article analyses 31 locked FY2025 public documents for a criterion-based maximum-variation sample of 24 listed EU/EEA firms. Four observed subgroups contain six firms each: energy, transport, digital infrastructure/electronic communications, and banking/financial-market infrastructure. The design combines deterministic candidate preparation with one human coder applying a frozen 0–3 codebook to 240 firm-domain cells. Purposeful variation supports structured comparison rather than population estimation (Patton 2015), and the documentary method remains limited to the meaning and provenance of the selected records (Bowen 2009).

The article makes three bounded contributions. Conceptually, it treats public-disclosure specificity as a distinct construct and makes its relationship to practice an explicit validity boundary (Adcock and Collier 2001). Empirically, it maps ten accountability, integration, assurance, continuity, and feedback domains without ranking firms. Methodologically, it links every code to evidence, preserves documented uncertainty, and tests non-exclusive bundles before considering whether a stable typology is justified. This last step follows the configurational insight that combinations can be informative while resisting the stronger claim that each case belongs to one latent organizational type (Fiss 2007).

The contribution is deliberately cumulative rather than classificatory. Domain distributions show where governed signals are more or less specific; transparent conjunction rules show which signals co-occur; and robustness tests show which conclusions change when documentary opportunity or uncertain codes change. This sequence permits recurring bundles to be described without converting them into mutually exclusive firm types. That restraint matters because typological claims require coherent, stable configurations and evidence that the proposed types add explanatory value beyond their component attributes (Doty and Glick 1994; Fiss 2007).

The maximum inference is intentionally narrow. The analyses identify what was disclosed in the reviewed corpus, with which level of specificity, and under which source and coding conditions. They do not establish European prevalence, sector effects, compliance, implementation quality, security, resilience in practice, or operational effectiveness. A code 0 means “not disclosed in the reviewed corpus”; a code 3 means that a governed signal was publicly specified. Neither value licenses a judgment about the underlying practice. That boundary controls the literature review, measures, results, and discussion that follow.

2. Literature review and conceptual framing

2.1 Corporate security governance as an accountability architecture

Corporate security governance can be approached as the allocation of authority, oversight, coordination, and assurance for security-relevant risks. In the information-security literature, governance distinguishes direction and accountability from the narrower execution of controls (Williams 2001). Frameworks accordingly connect board oversight, executive ownership, policy, risk management, and monitoring (Posthumus and von Solms 2004), while repeated management failures have been traced to treating security as a technical silo rather than an organizational responsibility (von Solms and von Solms 2004). The present study generalizes that accountability lens across corporate-security domains but does not assume that public statements faithfully describe internal arrangements.

The architecture has vertical and horizontal components. Vertically, boards and executives set direction, allocate ownership, and receive assurance. Horizontally, security functions coordinate with risk, internal control, audit, continuity, procurement, human resources, and operational units. Contemporary frameworks express comparable relationships: COSO integrates risk with strategy and performance (COSO 2017), and the NIST Cybersecurity Framework 2.0 places a Govern function around policy, roles, oversight, and supply-chain risk (National Institute of Standards and Technology 2024). These sources establish recognized governance interfaces; they do not provide a benchmark against which the sampled firms are certified or scored.

Auditability depends on the granularity of those public connections. A generic commitment to manage risk identifies an aspiration; a statement that links a named committee, accountable executive, reporting cadence, and escalation route identifies a governance pathway that an external reader can trace. Contemporary security-governance frameworks likewise distinguish direction, roles, policy, oversight, and monitoring (Williams 2001; National Institute of Standards and Technology 2024). The present construct captures that traceability in disclosure. It neither tests whether the pathway was used nor treats greater textual detail as proof of stronger control.

2.2 ERM integration, control, and assurance

ERM provides a public vocabulary for connecting security risks to enterprise risk registers, risk appetite, committees, internal control, and strategic decision processes. Empirical work shows that ERM implementation varies with organizational and governance conditions (Beasley, Clune, and Hermanson 2005), and interpretive work demonstrates that risk systems can embody different calculative cultures rather than one universal model (Mikes 2009). This contingency is important: integration may be named without revealing how it works in practice (Mikes and Kaplan 2015). ISO 31000 likewise supplies general principles for integrating risk management across organizational activities (International Organization for Standardization 2018). In this article, ERM integration is therefore coded as a disclosed interface, not as a proxy for ERM quality.

Assurance and feedback extend accountability beyond formal allocation. Publicly named testing, audit, certification, metrics, reporting cadence, escalation, corrective action, and lessons processes make it possible to observe whether a governance statement contains a control loop rather than a commitment alone. Risk specialists can also perform boundary work between expertise and managerial decision-making (Mikes 2011), which cautions against reducing integration to organizational charts. COSO’s emphasis on information, communication, review, and revision similarly demonstrates why feedback is part of governance (COSO 2017). The coding scheme treats these disclosed features as evidence of specificity and auditability; it cannot confirm the independence, coverage, frequency, or consequences of assurance activities.

2.3 Disclosure, signalling, and documentary specificity

Disclosure

theory helps explain why public records are analytically meaningful but incomplete. Signals can convey information when audiences cannot directly observe underlying qualities (Spence 1973), and corporate reporting can reduce information asymmetry under incentives and constraints (Healy and Palepu 2001; Verrecchia 2001). Public language can also support legitimacy by presenting organizational arrangements as appropriate to prevailing expectations (Suchman 1995). This study does not infer which motive produced any statement. It uses a more modest proposition: wording that identifies an owner, mechanism, process, recipient, metric, or feedback loop is more auditable as disclosure than a generic declaration, regardless of why it was published.

Risk-disclosure research provides direct comparators for that proposition. Annual-report studies document variation in narrative risk information (Linsley and Shrives 2006), associate disclosure patterns with corporate characteristics (Abraham and Cox 2007), and use content analysis to distinguish types of risk statements (Deumes 2008). Readability research further shows that textual properties affect the information environment and should not be treated as neutral packaging (Loughran and McDonald 2014). The present scale is not a readability or quantity measure. It focuses on the governance content of a qualifying passage and separates generic acknowledgement from a named mechanism and from a governed signal with accountability or feedback.

Selectivity also limits comparisons based on silence. Firms face different reporting conventions, materiality judgements, document architectures, and incentives to condense or elaborate particular topics. Narrative form can itself affect what automated or human readers recover from reports (Verrecchia 2001; Loughran and McDonald 2014). A missing signal can therefore arise from non-disclosure, dispersion across documents, or an opportunity boundary in the collected corpus. The zero rule used here responds by assigning a precise documentary meaning—no qualifying disclosure located—without extending that meaning to an absent organizational practice.

2.4 Critical-service context and third-party dependence

Critical-service firms operate in regulatory and dependency environments that make governance publicly salient. NIS2 establishes cybersecurity risk-management and reporting expectations across essential and important sectors, including governance and supply-chain security (European Parliament and Council 2022a). The Critical Entities Resilience Directive addresses the ability of designated entities to prevent, resist, mitigate, respond to, and recover from disruptive incidents (European Parliament and Council 2022b). These instruments justify observing several sectors through a shared governance lens. They do not make the sample a compliance audit, and the study does not determine whether any firm falls within a particular national designation or meets a legal obligation.

Third-party governance deserves separate attention because critical services depend on suppliers, contractors, technology providers, and inter-organizational infrastructures. Cyber supply-chain scholarship emphasizes reduced visibility and the strategic control of distributed dependencies (Boyson 2014). NIST guidance translates the issue into policies, roles, due diligence, risk assessment, monitoring, and response across organizational levels (Boyens et al. 2024). Public evidence may nevertheless be dispersed across procurement, cyber, sustainability, and contractual documents. A low third-party code in a fixed source set therefore has a particularly strong non-inference rule: it identifies what was not located, not what the firm failed to do.

A third-party signal becomes more auditable when it moves beyond a broad reference to supplier risk. Named due-diligence criteria, contractual security obligations, monitoring, incident notification, assurance rights, and remediation routes each expose a different part of the governance chain. Supply-chain scholarship and guidance emphasize that dependencies cross organizational boundaries and cannot be reduced to internal technical controls (Boyson 2014; Boyens et al. 2024). Coding these named mechanisms separately from generic risk language makes the documentary threshold explicit while leaving their operational reach untested.

2.5 From documents to bounded measures

Moving from documents to an ordinal measure requires construct discipline. Validity depends on the relationship between the systematized concept and the observations used to represent it (Adcock and Collier 2001). Content analysis therefore requires explicit units, a coding frame, decision rules, and an audit trail (Krippendorff 2018; Neuendorf 2017; Schreier 2012). The 0–3 scale used here is ordered by public specificity: no qualifying disclosure, generic commitment, named mechanism, and governed signal. It is not an interval scale, does not aggregate naturally into corporate performance, and is interpreted only within the locked source hierarchy.

Cross-domain combinations create a second measurement choice. Typologies make strong claims: they posit coherent configurations and often assign cases to distinct types (Doty and Glick 1994). Set-theoretic approaches instead draw attention to combinations and equifinality (Fiss 2007). With 24 cases and an ordinal disclosure measure, transparent rule-based bundles are the more defensible first step. They show which domains co-occur above an explicit threshold and permit overlap. A named mutually exclusive typology is retained only if assignments are substantively coherent and stable under alternative thresholds, uncertainty, and evidence-scope restrictions.

2.6 Analytical expectations

The framework yields analytical expectations rather than effectiveness hypotheses. First, domains embedded in standardized corporate risk and control vocabularies may be disclosed more specifically than mechanisms distributed across specialist documents. Second, recurring domain bundles may be visible without forming exclusive firm types. Third, observed subgroup profiles may differ descriptively while remaining too uncertain for population inference. Fourth, restricting the evidence base or applying documented alternative codes may narrow the findings. Each expectation is evaluated against denominators and pre-specified sensitivities, with adverse results retained. None predicts security outcomes, compliance, or the quality of internal implementation.

3. Materials and methods

3.1 Design and inferential boundary

The study used a structured public-document design to compare the specificity and audibility of disclosed corporate security-governance signals. Documents were treated as social and organizational records whose content could be systematically located and interpreted (Bowen 2009), not as transparent windows onto practice. The analytical unit was the firm-year/domain cell, supported by document- and passage-level provenance. The design followed content-analysis principles of explicit units, categories, coding instructions, and reproducible transformations (Krippendorff 2018). The inferential target was the locked FY2025 corpus; no claim was made about unobserved documents, internal operations, European firm populations, or causal effects.

3.2 Sample and documentary opportunity

The criterion-based maximum-variation sample contained 24 publicly listed EU/EEA firms, divided into four observed critical-service subgroups of six: energy; transport; digital infrastructure/electronic communications; and banking/financial-market infrastructure. Selection sought cross-sector and cross-jurisdictional variation while retaining a common reporting period and public-document basis. Purposeful samples select information-rich cases for analytic comparison rather than probabilistic representation (Patton 2015). Case diversity can strengthen examination of mechanisms and boundaries (Flyvbjerg

2006), but it does not create sampling weights or authorize prevalence estimates. Table 1 identifies only neutral firm IDs, countries, subgroups, and source coverage.

The source hierarchy required a controlling FY2025 annual or integrated report for every firm. Supplementary corporate-governance, sustainability/ESG, or public security/cyber/resilience documents were eligible when they met the pre-specified authority, period, accessibility, and relevance rules. Retrieval, deduplication, source-class decisions, and access details were frozen before analysis. The final lock contained 31 documents: 19 firms contributed one source, three contributed two, and two contributed three. Because every multi-source firm was in the energy subgroup, documentary opportunity was recorded as a design condition and addressed through sensitivity analysis, not treated as disclosure quality.

3.3 Domains, unit, and ordinal specificity scale

Ten domains operationalized the governance architecture: board oversight (D01), executive accountability (D02), security-function architecture (D03), ERM integration (D04), incident escalation and learning (D05), assurance and audit (D06), workforce and culture (D07), third-party governance (D08), continuity and resilience (D09), and metrics and reporting (D10). The domains were informed by governance, enterprise-risk, and cybersecurity frameworks (Williams 2001; COSO 2017; National Institute of Standards and Technology 2024) but were applied as disclosure categories rather than compliance criteria. Twenty-four firms by ten domains produced 240 primary cells.

Each cell received one human-verified ordinal code. Code 0 meant no qualifying signal was disclosed in the reviewed source set after the required search. Code 1 captured a generic commitment or acknowledgement without a named mechanism. Code 2 required a named mechanism, role, process, requirement, or structured arrangement. Code 3 required a governed signal with identifiable accountability, cadence, assurance, metric, escalation, monitoring, or feedback. The scale ordered documentary specificity, not substantive performance. This construct-to-indicator rule was fixed to protect measurement validity (Adcock and Collier 2001), and examples and edge cases were versioned in the codebook (Schreier 2012).

The scale was applied as an ordered evidentiary rubric, not as an interval measure of governance quality. Movement from 0 to 1 records the appearance of qualifying but generic language; movement to 2 requires a named mechanism or accountable actor; and movement to 3 requires a governed connection that is sufficiently explicit to audit in the source. Because ordinal categories express progressively specific documentary criteria, arithmetic distances between codes were not interpreted as equal improvements (Krippendorff 2018; Schreier 2012). Threshold counts and complete distributions therefore remain the principal summaries.

3.4 Evidence preparation and human coding

A deterministic preparation stage converted eligible documents to searchable text, executed domain-specific queries, and assembled candidate passages with source and location fields. This stage organized evidence but did not assign interpretive 0–3 codes. One human coder then reviewed candidates, reopened source context as required, applied the frozen codebook, selected supporting evidence, and completed the zero-search record for code-0 decisions. Separating mechanical retrieval from interpretive judgment preserves an auditable sequence without representing software output as human coding. Codebook stability, explicit units, and decision documentation follow established content-analysis practice (Neuendorf 2017; Krippendorff 2018).

The project used one human coder; it did not use two independent coders, a C1/C2 split, or an inter-rater coefficient. Reliability controls were therefore designed around transparency and within-coder consistency. The workflow included a pilot, human freeze of the codebook, full review of flagged or ambiguous cases, evidence locators, decision notes, and a delayed intra-coder control. Coding-unit definition and decision consistency are recognized sources of difficulty even in multi-coder designs (Campbell et al. 2013). The present approach reports its single-coder limitation directly and makes the documented uncertainties available to sensitivity analysis rather than implying independent agreement (Schreier 2012).

Twenty-eight cells retained a documented alternative code after human review. These were not silently resolved in the direction of the principal claim. For each cell, the register preserved the primary code, alternative code, source status, delta, and whether the alternative crossed the ≥ 2 or $= 3$ thresholds. The primary dataset was frozen after human verification. Sensitivity R3 then examined uncertainty exclusion and the simultaneous application of all documented alternatives. This design makes conclusion changes visible while avoiding a post hoc search for individually favorable recodings.

The audit trail treated ambiguity as data rather than as a nuisance to erase. For every uncertain cell, the register retained the primary code, the plausible alternative, the evidence locator, and the decision basis. A delayed intra-coder review checked the entire coded set against the versioned codebook, while the uncertainty sensitivity replaced all documented alternatives together. This procedure improves transparency but does not substitute for independent coding or an inter-rater estimate (Campbell et al. 2013; Neuendorf 2017). The resulting 28 uncertain cells therefore remain an explicit limitation.

3.5 Primary analyses

Primary analyses were pre-specified before execution. Sample and source coverage were summarized first. Domain distributions then reported counts at codes 0–3, evaluable denominators, medians, ranges, and thresholds ≥ 1 , ≥ 2 , and $= 3$. The firm-domain matrix was visualized in an order fixed solely by subgroup and neutral firm ID. No row total, mean, composite, rank, leader/laggard label, or best/worst comparison was published. The relative positions of D04 and D08 were reported only as within-corpus patterns of disclosure specificity and were linked to the robustness tests that could change those conclusions.

Five transparent rules assessed cross-domain combinations. Governance core required D01–D04 all ≥ 2 ; assurance-feedback required D06 and D10 both ≥ 2 ; response-resilience required D05 and D09 both ≥ 2 ; the broad profile required all ten domains ≥ 2 ; and the third-party disclosure gap required D08 ≤ 1 . Rules were non-exclusive, so one firm could meet several bundles. Candidate mutually exclusive labels were examined but could be retained only if assignments remained coherent under threshold and evidence-scope changes. This distinguishes descriptive combinations from stronger typological claims (Doty and Glick 1994; Fiss 2007).

3.6 Observed subgroup contrasts

For the four observed subgroups, the analysis displayed domain medians, ranges, coverage, and uncertainty counts. Each subgroup contained six purposively selected firms. Domain-level differences were assessed with 20,000 label permutations using the pre-specified statistic, an approach suited to randomization-based inference under the stated exchangeability design (Good 2005). Ten resulting p-values were adjusted with Holm's sequentially rejective procedure (Holm 1979). Because the sample was not probabilistic and documentary opportunity was uneven, the tests were interpreted as bounded diagnostics. They did not estimate sector populations or identify sector effects.

3.7 Robustness and conclusion-change controls

Seven robustness analyses were executed. R1 restricted support to annual or integrated reports and treated positive cells supported only by supplementary documents as not evaluated, never as zero. R2 compared ordinal, ≥ 2 , and $= 3$ readings. R3 applied uncertainty exclusions and documented alternatives. R4 repeated central domain conclusions after removing each firm in turn. R5 used 50,000 permutations to diagnose the association between source count and an unweighted disclosure-specificity summary, with the summary prohibited as a firm-ranking result. R6 deterministically excluded evidence explicitly limited to cyber/information security without recoding cells. R7 confirmed that the unweighted composite remained diagnostic only.

Each robustness test carried a pre-specified conclusion-change rule. A supported result could be reported in its bounded form; a qualified result required the relevant denominator, range, scope restriction, or non-inference in the same paragraph; and an unsupported result was suppressed. Claims about a mutually exclusive typology and a positive document-density association were not retained as findings. All results unfavourable to the candidate typology, subgroup interpretation, source-density association, or balanced corporate-security breadth were retained.

Before drafting, each planned claim was assigned a controlling analysis, exhibit, denominator, sensitivity requirement, and permitted wording. Claims were reported only within those pre-specified limits. In particular, the named typology and any positive document-density association were not presented as findings.

3.8 Reproducibility and input integrity

Reproducibility controls separated immutable inputs from derived outputs. The frozen analytical dataset, codebook, analysis plan, and notebook received versioned filenames and SHA-256 hashes. A run-all execution generated tables, figures, statistical outputs, and a scientific checksum. The anonymous reviewer repository contains a privacy-scrubbed analytical copy, the executable notebook, and package-specific hashes. These identifiers control input identity; they do not imply external validation.

3.9 Ethics, rights, and availability boundary

The study analysed public corporate documents and did not recruit participants, intervene, or collect non-public personal data. Corporate PDFs were used as sources but were not treated as proof of effectiveness and are not redistributed. Firm identities are represented by neutral IDs in the analytical exhibits. An anonymized analytical dataset, codebook, source manifest, analysis notebook, and reproducibility outputs are available through the anonymous OSF view-only link stated in the Data and code availability declaration.

4. Results

4.1 Sample and documentary opportunity

The reviewed corpus comprised 24 publicly listed critical-service firms. The criterion-based maximum-variation design allocated six firms to each of four observed subgroups: Energy, Transport, Digital infrastructure/electronic communications, and Banking/financial market infrastructure. The source lock contained 31 corporate documents. Every firm contributed a controlling annual or integrated report; 19 firms had one locked source, three had two, and two had three. These counts describe the sampled corpus and do not support a claim of representativeness for European firms or sector populations (Table 1).

Table 1
Sample sectors and source coverage

firm id	sector	country	source count	document classes	annual report available	locked source coverage	coded domains n	uncertainty cells n
F01	Energy	France	1	AR	True	COMPLETE	10	1
F02	Energy	Italy	2	AR CG	True	COMPLETE	10	1
F03	Energy	Spain	3	AR CG ESG	True	COMPLETE	10	1
F04	Energy	Germany	3	AR CG POL	True	COMPLETE	10	2
F05	Energy	Finland	2	AR CG	True	COMPLETE	10	3
F06	Energy	Portugal	2	AR CG	True	COMPLETE	10	0
F07	Transport	Germany	1	AR	True	COMPLETE	10	1
F08	Transport	Spain	1	AR	True	COMPLETE	10	0
F09	Transport	Denmark	1	AR	True	COMPLETE	10	4
F10	Transport	France	1	AR	True	COMPLETE	10	1
F11	Transport	Austria	1	AR	True	COMPLETE	10	3
F12	Transport	Estonia	1	AR	True	COMPLETE	10	2
F13	Digital infrastructure / electronic communications	Germany	1	AR	True	COMPLETE	10	2
F14	Digital infrastructure / electronic communications	France	1	AR	True	COMPLETE	10	0
F15	Digital infrastructure / electronic communications	Spain	1	AR	True	COMPLETE	10	1
F16	Digital infrastructure / electronic communications	Norway	1	AR	True	COMPLETE	10	0
F17	Digital infrastructure / electronic communications	Netherlands	1	AR	True	COMPLETE	10	0
F18	Digital infrastructure / electronic communications	Belgium	1	AR	True	COMPLETE	10	1
F19	Banking / financial market infrastructure	France	1	AR	True	COMPLETE	10	0
F20	Banking / financial market infrastructure	Netherlands	1	AR	True	COMPLETE	10	1
F21	Banking / financial market infrastructure	Finland	1	AR	True	COMPLETE	10	3
F22	Banking / financial market infrastructure	Spain	1	AR	True	COMPLETE	10	0
F23	Banking / financial market infrastructure	Netherlands	1	AR	True	COMPLETE	10	1
F24	Banking / financial market infrastructure	Germany	1	AR	True	COMPLETE	10	0
<i>Note. Codes describe public-disclosure specificity only; code 0 means not disclosed in the reviewed corpus.</i>								

Documentary opportunity was uneven across the four observed subgroups. All five firms with more than one locked source were in Energy, whereas the other three subgroups were represented by annual or integrated reports alone. Accordingly, source count was treated as an opportunity condition and a design limitation, not as an indicator of disclosure quality or corporate practice. The source-density diagnostic was also sector-confounded and is reported only as a sensitivity result (Table 1; Table A1).

4.2 Domain distributions

Across the 240 firm-domain cells, the observed code distribution was 12 at code 0, 14 at code 1, 30 at code 2, and 184 at code 3; 228 cells therefore contained a qualifying public-disclosure signal. These values summarize the specificity and auditability of disclosure in the reviewed source set. In particular, code 0 means only that a qualifying signal was not disclosed in the reviewed corpus; it does not demonstrate that the corresponding practice was absent, deficient, non-compliant or ineffective (Table 2).

ERM integration (D04) was the most consistently disclosed governed signal in the reviewed corpus: 23 of 24 firms were coded 3, and all 24 reached the named-mechanism threshold of at least 2. This observation concerns the visibility of integration into enterprise risk and control processes, not the quality or

effectiveness of ERM in practice. The D04 position remained stable across the relevant source-scope, threshold, uncertainty and leave-one-firm-out checks (Table 2; Table A1).

Third-party governance (D08) was least often disclosed at the named-mechanism threshold: 16 of 24 firms reached at least code 2, while five of 24 were coded 0. This is an asymmetry in public-disclosure specificity. The five zeros mean not disclosed in the reviewed corpus and cannot be read as evidence that supplier or third-party controls were absent. D08 remained among the lowest domains at the mechanism threshold throughout the relevant robustness checks (Table 2; Table A1).

The remaining eight domains also contained many code-3 signals, but their distributions were not collapsed into a global score. Table 2 retains the four-level counts, denominators, medians and ranges for every domain so that generic commitments, named mechanisms and governed signals remain distinguishable. The table is therefore a profile of disclosure specificity, not a corporate-security performance scale.

Table 2
Ten-domain definitions and observed distributions

domain id	domain	definition short	evaluable n	code 0 n	code 0 pct	code 1 n	code 1 pct	code 2 n	code 2 pct	code 3 n	code 3 pct	median	min	max
D01	Board oversight	Board-level allocation and recurring security-related oversight.	24	0	0.0%	1	4.2%	2	8.3%	21	87.5%	3.0	1.0	3.0
D02	Executive accountability	Named executive ownership, remit, and reporting.	24	0	0.0%	2	8.3%	4	16.7%	18	75.0%	3.0	1.0	3.0
D03	Security-function architecture	Named security function, programme, and coordination architecture.	24	0	0.0%	3	12.5%	2	8.3%	19	79.2%	3.0	1.0	3.0
D04	ERM integration	Integration into enterprise risk and control processes.	24	0	0.0%	0	0.0%	1	4.2%	23	95.8%	3.0	2.0	3.0
D05	Incident escalation and learning	Incident escalation, corrective action, and learning.	24	1	4.2%	1	4.2%	6	25.0%	16	66.7%	3.0	0.0	3.0
D06	Assurance and audit	Security-scoped audit, testing, certification, or assurance.	24	2	8.3%	1	4.2%	2	8.3%	19	79.2%	3.0	0.0	3.0
D07	Workforce and culture	Security training, awareness, culture, and monitored participation.	24	2	8.3%	1	4.2%	3	12.5%	18	75.0%	3.0	0.0	3.0
D08	Third-party governance	Security-scoped supplier and third-party requirements and oversight.	24	5	20.8%	3	12.5%	5	20.8%	11	45.8%	2.0	0.0	3.0
D09	Continuity and resilience	Governed continuity, recovery, crisis, or resilience processes.	24	2	8.3%	1	4.2%	2	8.3%	19	79.2%	3.0	0.0	3.0
D10	Metrics and reporting	Security metrics, targets, trends, cadence, and reporting recipients.	24	0	0.0%	1	4.2%	3	12.5%	20	83.3%	3.0	1.0	3.0

Note. Codes describe public-disclosure specificity only; code 0 means not disclosed in the reviewed corpus.

4.3 Firm × domain profiles

Figure 1 displays the 24 by 10 firm-domain matrix in a fixed order defined only by sector and then neutral firm identifier. The heatmap shows that code-3 signals are widespread across several governance domains while lower-specificity cells are distributed across firms and are visibly more concentrated in D08. Because rows are not sorted by a sum, mean or composite, the visual does not produce or imply a firm ranking (Fig. 1).

Within a firm row, the profile represents what was located under the locked source and coding rules rather than a complete inventory of internal governance. A lighter cell therefore signals less specific public disclosure in the reviewed corpus, including the possibility of non-disclosure, but not a weaker underlying practice. Figure 1 should be read across domains and patterns, not as a league table.

4.4 Non-exclusive disclosure bundles

Three transparent, rule-based and non-exclusive bundles appeared frequently at the named-mechanism threshold. The governance core (D01-D04 all at least 2) was present in 19 of 24 firms; the assurance-feedback bundle (D06 and D10 both at least 2) in 20 of 24; and the response-resilience bundle (D05 and D09 both at least 2) in 21 of 24. Because the rules overlap, a firm can satisfy more than one bundle and the counts do not define mutually exclusive firm types (Table 3).

A broad profile in which all ten domains reached at least code 2 was observed for 14 of 24 firms under the primary codes and for 13 of 24 when all documented alternative codes were applied simultaneously. The defensible result is therefore a range of 13–14 firms, not a fixed maturity category. Separately, eight of 24 firms met the pre-specified third-party disclosure-gap rule (D08 at code 0 or 1), again as a non-exclusive disclosure bundle rather than evidence of deficient third-party practice (Table 3; Tables A1-A2).

The candidate labels ownership-led, assurance-led, resilience-led and fragmented were not retained as stable, mutually exclusive types. Category allocation changed materially when the threshold moved from at least 2 to code 3, and the evidence-scope sensitivity sharply reduced complete profiles. The analysis therefore retains only the transparent overlapping bundles and reports the non-retention of the named typology as an adverse result (Table 3; Table A1).

Table 3
Rule-based non-exclusive configuration assessment

configuration id	candidate label	rule	primary support n	primary evaluable n	primary support pct	annual only support n	annual only evaluable n	alternative support n	alternative evaluable n	uncertainty excluded support n	uncertainty excluded evaluable n
CFG-01	Governance core > = 2	D01-D04 all > = 2	19	24	79.2%	16	21	21	24	16	16
CFG-02	Assurance-feedback bundle > = 2	D06 and D10 both > = 2	20	24	83.3%	19	23	19	24	17	17
CFG-03	Response-resilience bundle > = 2	D05 and D09 both > = 2	21	24	87.5%	20	23	19	24	18	21
CFG-04	Broad ten-domain profile > = 2	D01-D10 all > = 2	14	24	58.3%	13	21	13	24	5	8
CFG-05	Third-party disclosure gap	D08 < = 1	8	24	33.3%	8	23	7	24	7	18

Note. Codes describe public-disclosure specificity only; code 0 means not disclosed in the reviewed corpus.

4.5 Observed subgroup contrasts

Figure 2 summarizes medians, observed ranges, uncertainty counts and source coverage for four subgroups of six firms each. Median profiles were compressed near code 3 for several domains. D08 showed the clearest descriptive separation: the subgroup medians were 1.5 in Energy, 1 in Transport, 2 in Digital infrastructure, and 3 in Banking/financial market infrastructure. The first three subgroups nevertheless each spanned codes 0–3, while the six Banking/financial market infrastructure observations were all code 3. These are within-sample descriptions only (Fig. 2; n = 6 per subgroup).

For each domain, the subgroup comparison used 20,000 label permutations and Holm adjustment across the ten tests. None of the ten adjusted p-values was below 0.05; for D08, the raw permutation p-value of 0.0308 became 0.3085 after adjustment. The observed contrasts therefore remain descriptive and cannot be generalized to sector populations or interpreted as sector effects (Fig. 2; Table A1; n = 6 per subgroup).

4.6 Robustness and conclusion-change controls

R1 restricted the evidence basis to annual or integrated reports. Of the 240 cells, 230 remained evaluable; ten positive cells supported only by non-annual documents were treated as not evaluated, never as code 0. The central ordering of D04 and D08 remained stable under this conservative source restriction. R4 reached the same ordering in all 24 leave-one-firm-out runs, showing that the within-sample observation was not driven by one firm (Table A1).

R2 preserved the D04/D08 domain ordering across the ordinal and binary readings but weakened the candidate named typology. R3 examined all 28 uncertainty cells through exclusion and documented alternative codes. The broad ten-domain profile changed from 14 of 24 to 13 of 24 under the simultaneous alternatives, which is why the main text reports a 13–14 range and retains the adverse sensitivity rather than selecting the more favorable count (Tables A1-A2).

R5 used the unweighted disclosure-specificity summary solely as a diagnostic of document opportunity. The association with source count was $\rho = -0.3012$, with a two-sided 50,000-permutation p-value of 0.1556. Because every multi-source firm belonged to Energy, the result was classified INCONCLUSIVE / SECTOR_CONFOUNDED. It does not support an association claim and instead limits the interpretation of subgroup contrasts (Table A1).

R6 classified 149 of the 228 non-zero evidence records as CYBER_ONLY_EXPLICIT without recoding any cell. After those records were excluded, only 91 of 240 cells remained evaluable and only one of 24 firms retained a complete ten-domain profile. This is a material scope restriction: the observed governance evidence is predominantly mediated by cyber/information-security disclosure and cannot be presented as balanced coverage of all corporate-security functions (Table A1).

R7 confirmed that the unweighted composite had a diagnostic role only. No firm-level composite table, order or ranking was produced, and the composite is not treated as a primary result. Its sole reported function is to document the inconclusive R5 opportunity check (Table A1).

5. Discussion

5.1 What public disclosure makes visible

The findings show that public reporting can make substantial parts of a security-governance architecture visible: named oversight bodies, executive ownership, enterprise-risk interfaces, assurance mechanisms, continuity arrangements and reporting loops were frequently disclosed in the sampled corpus. That visibility is analytically useful because it can be traced to specific documents and coded against consistent criteria. It remains a property of disclosure, however, and cannot validate whether the named arrangements operate as described or produce effective security outcomes. (Healy and Palepu 2001; Linsley and Shrivs 2006; Deumes 2008)

A signal-based reading should remain equally cautious about motive. Specific language may inform stakeholders, respond to regulation, demonstrate legitimacy, or standardize internal narratives (Spence 1973; Verrecchia 2001; Suchman 1995). The coding architecture does not choose among these explanations. Its value is comparative and procedural: a reader can locate the passage, see why it crossed a threshold, and identify which governance element was publicly named. That is stronger than counting security-related words, but still weaker than observing implementation.

5.2 The D04–D08 asymmetry

The clearest substantive pattern is the asymmetry between ERM integration (D04) and third-party governance (D08). D04 was almost uniformly disclosed as a governed signal, whereas D08 was less consistently visible at the named-mechanism threshold and contained more non-disclosures. The contrast is best understood as an asymmetry in public specificity and auditability, not as a difference in internal capability. (Mikes 2009; Boyson 2014)

One plausible interpretation is that enterprise-risk structures are more standardized and more readily expressed through established annual-report vocabularies, such as risk committees, internal controls and risk appetite. The documentary design cannot test that explanation or infer reporting motives. It can establish only that D04 was more consistently articulated under the coding rules and that this ordering was not driven by one sampled firm. (COSO 2017; Mikes and Kaplan 2015)

Conversely, third-party governance may be distributed across procurement, sustainability, cyber-risk and contractual materials that are not equally visible in a fixed public source set. That possibility reinforces, rather than relaxes, the zero rule: lower D08 codes identify a disclosure boundary in this corpus. They do not establish that supplier due diligence, contractual controls or monitoring were absent. (Boyson 2014; Boyens et al. 2024)

The asymmetry is particularly relevant because governance frameworks and NIS2 make supply-chain risk an explicit organizational concern (National Institute of Standards and Technology 2024; Boyens et al. 2024; European Parliament and Council 2022a). The present evidence cannot determine whether lower D08 specificity reflects dispersed reporting, commercial sensitivity, less standardized vocabulary, or another cause. It therefore supports a research question about disclosure architecture, not a diagnosis of supplier controls. The 24/24 leave-one-out stability shows only that the D04–D08 ordering is not attributable to one sampled firm.

5.3 Non-exclusive bundles rather than stable firm types

The overlapping governance, assurance-feedback and response-resilience bundles offer a transparent way to summarize recurring combinations without forcing each firm into a single category. Their rules are inspectable and preserve the domain profile as the primary object. The broad-profile range of 13–14 firms also demonstrates why uncertainty should be reported as a range when documented alternative codes cross the selected threshold. (Doty and Glick 1994; Fiss 2007)

The non-retention of the candidate named typology is itself informative. A stable type should not depend materially on whether a named mechanism or a governed signal is used as the threshold, nor should it collapse when the dominant evidence vocabulary is restricted. Here, both problems occurred. Reporting the adverse assessment prevents descriptive bundles from being mistaken for latent organizational types. (Doty and Glick 1994; Fiss 2007)

This negative typology result also clarifies the level at which configuration language is warranted. Configurational theory can accommodate multiple paths and combined attributes (Fiss 2007), but a typology ordinarily asserts more than co-occurrence: it proposes coherent ideal types with explanatory content

(Doty and Glick 1994). Here, threshold dependence and the cyber-scope restriction remove that warrant. The bundles remain useful descriptive summaries because their rules and denominators are transparent, overlapping, and reversible.

5.4 Bounded observed-subgroup contrasts

The subgroup profiles suggest contextual variation worth investigating, particularly for D08, but the present design does not estimate sector effects. Each subgroup contains only six purposively selected firms, none of the ten Holm-adjusted tests crossed 0.05, and documentary opportunity was uneven and sector-confounded. Figure 2 should therefore be read as a map of four observed subgroups, not as a comparison of European sector populations. (Patton 2015; Flyvbjerg 2006)

5.5 Cyber-mediated evidence and the breadth of corporate security

The R6 sensitivity materially narrows the breadth of the evidence base. Almost two thirds of non-zero records were explicitly cyber-only, and excluding them left too little complete coverage for a balanced ten-domain comparison. The article can still examine corporate security governance as publicly disclosed, but it must state that this public articulation is predominantly mediated by cyber and information-security language. (Williams 2001; National Institute of Standards and Technology 2024)

The cyber mediation is understandable because information-security governance has developed explicit vocabularies for roles, oversight, controls, assurance, and metrics (Williams 2001; Posthumus and von Solms 2004; National Institute of Standards and Technology 2024). It nevertheless creates construct imbalance. Physical protection, personnel security, fraud, crisis management, and other corporate-security functions may be less visible in annual reports or expressed outside cyber terminology. The article therefore reports a corporate-security governance profile whose public evidence is predominantly cyber-mediated, not balanced coverage of the broader field.

5.6 Methodological contribution: disclosure is not practice

The principal methodological contribution is the separation of public-disclosure specificity from practice. The ordinal rule asks whether a public source moves from a generic commitment to a named mechanism and then to a governed signal with assurance, metrics or feedback. That sequence improves auditability of documentary comparison, but it is not a measure of implementation, compliance, resilience or operational effectiveness. (Adcock and Collier 2001; Bowen 2009)

The evidence-linked ten-domain profile also preserves distinctions that a single composite would obscure. Domain distributions, overlapping bundles and conclusion-change tests can be audited separately, while firm ordering is deliberately withheld. This structure is reusable for later studies provided that source opportunity, coding scope and the disclosure-practice boundary are re-established rather than assumed. (Krippendorff 2018; Neuendorf 2017)

The conclusion-change architecture is as important as the scale itself. Construct validity requires evidence that the operational rule behaves in ways consistent with the intended concept (Adcock and Collier 2001). By retaining alternative codes, source-scope exclusions, threshold changes, and adverse configuration results, the analysis shows where the public-disclosure construct is stable and where it is not. This audit trail is consistent with systematic content analysis (Krippendorff 2018; Schreier 2012) and prevents one favorable reading from being mistaken for a definitive classification.

Testing implementation would require a different evidentiary design. Matched internal policies, committee records, assurance work, incident data, interviews, and longitudinal observations could be triangulated with public statements to examine correspondence and divergence (Jick 1979; Bowen 2009). Even then, outcome attribution would require careful control of exposure and context. The present public-document design supplies neither those internal observations nor a counterfactual. Its value is to make the external evidence boundary reproducible, so that later studies can test rather than assume the relationship between disclosure and practice.

5.7 Practical relevance as a disclosure audit

For boards and security leaders, the ten domains may serve as a neutral prompt for reviewing what an organization makes publicly visible about ownership, integration, assurance and feedback. Such use would be a disclosure audit, not an assessment of whether the underlying controls are sufficient or effective. The present results do not justify prescriptive thresholds or comparative performance judgments. (Williams 2001; COSO 2017)

5.8 Limitations

First, the sample was purposive and limited to 24 listed EU/EEA firms across four critical-service subgroups. It was designed for structured variation, not statistical representation. Second, the analysis was bounded to 31 locked public documents for FY2025 reporting, and source opportunity differed across subgroups. The annual-report restriction mitigated but did not remove the possibility that relevant mechanisms were disclosed elsewhere.

Third, coding was completed by one human coder using a version-controlled codebook, evidence locators, full review of flagged cases and a delayed intra-coder control. This design emphasizes transparent decisions rather than intercoder agreement. The 28 uncertainty cells were retained and tested; their documented alternatives changed the broad-profile count by one firm and are therefore visible in the reported range. (Schreier 2012; Campbell et al. 2013)

The single-coder design remains a material limitation even with these controls. Unitization and interpretive consistency are recurring coding challenges (Campbell et al. 2013), and an independent coder could reveal disagreements that a delayed intra-coder check cannot. The study's response is transparency rather than equivalence: evidence locators, frozen definitions, uncertainty alternatives, and conclusion-change tests expose decisions for audit (Neuendorf 2017). Replication should assess whether other trained coders reproduce the same domain patterns without retrospectively changing the present design.

Fourth, the evidence base was heavily mediated by cyber/information-security disclosure, which limits claims about corporate security as a broader field. Finally, public documents provide no direct access to internal implementation or operational outcomes. Neither a code 0 nor a code 3 can establish absence,

quality or effectiveness of practice; both describe the public evidence available under the study rules.

5.9 Research agenda

Future work could test whether the disclosure patterns persist in larger pre-specified samples, other reporting years and additional jurisdictions. Triangulation with interviews, internal governance records or regulatory filings could examine the relationship between public specificity and practice, while independent replications could test the stability of the ten-domain codebook. Those designs would address questions that the present public-document analysis cannot answer. (Jick 1979; Bowen 2009)

A second research stream should connect disclosure to practice without collapsing the two constructs. Triangulation can combine documents with interviews, internal records, observations, or regulatory evidence while preserving the different claims each source supports (Jick 1979; Bowen 2009). A longitudinal design could also test whether public specificity changes after regulation, incidents, or governance reorganization. Such work would require new access, ethics, and causal design decisions. It cannot be inferred from the present cross-sectional public corpus.

Replication should also vary the cases and the documentary window. Probability-based prevalence claims would require a defensible population frame, while theory-driven case extensions could test whether the same domain asymmetry appears in other jurisdictions, ownership structures, or reporting regimes (Patton 2015; Flyvbjerg 2006). Repeated annual coding could separate changes in disclosure practice from one-period document availability. These extensions should preserve denominators, neutral firm identifiers, uncertainty alternatives, and conclusion-change thresholds so that apparent longitudinal or cross-context differences remain auditable.

5.10 Discussion synthesis

Taken together, the analysis supports a restricted conclusion: the sampled public disclosures make some governance signals, especially ERM integration, highly visible; third-party governance is less consistently specific; and overlapping bundles recur without supporting a stable mutually exclusive typology. These within-sample observations survive the relevant central checks, but their breadth is constrained by cyber-mediated evidence and they do not measure organizational practice.

6. Conclusion

This study asked what 24 listed EU/EEA critical-service firms made publicly visible about the allocation, integration, and assurance of corporate security governance. Within the 31-document FY2025 corpus, ERM integration was the most consistently disclosed governed signal: 23/24 firms were coded 3 and 24/24 reached the named-mechanism threshold. Third-party governance was least often disclosed at that threshold: 16/24 reached ≥ 2 and five cells were code 0. The ordering remained stable in all 24 leave-one-firm-out analyses. These are within-corpus findings about public specificity; they do not establish the quality, presence, or effectiveness of the corresponding practices.

Three transparent non-exclusive bundles recurred, while the broad ten-domain profile applied to 13–14 firms depending on documented uncertainty alternatives. The candidate mutually exclusive typology was not retained because assignments were threshold-sensitive and evidence-scope restriction sharply reduced complete profiles. Reporting that adverse result is part of the contribution: recurring combinations can be described without converting them into stable organizational types or rankings.

The principal methodological contribution is an evidence-linked separation of disclosure specificity from practice. Explicit domains, source locators, the zero rule, uncertainty alternatives, and conclusion-change tests make the documentary comparison auditable (Adcock and Collier 2001; Krippendorff 2018). The boundary is also substantive: 149/228 non-zero evidence records were explicitly cyber-only, and their exclusion left 91/240 evaluable cells and one complete profile. The article therefore characterizes the evidence as predominantly cyber/information-security mediated.

A code 0 remains “not disclosed in the reviewed corpus,” and a code 3 remains a governed public signal; neither validates operational reality. The purposive sample, uneven documentary opportunity, single human coder, 28 uncertainty cells, and lack of internal evidence rule out claims about European prevalence, sector effects, compliance, resilience in practice, or effectiveness. Future replication and mixed-source research can test those relationships. The present article offers a controlled way to state what public documents make visible—and to stop precisely where those documents stop.

Declarations

Competing interests

The author has no relevant financial or non-financial interests to disclose.

Consent to participate

Not applicable; the study did not involve human participants.

Consent for publication

Not applicable; no identifiable participant data are reported.

Use of generative AI

Generative artificial intelligence tools were used to support project structuring, data-processing code, analytical checking, and language development. The author selected the sources, verified the evidence and references, reviewed the code and outputs, made all scientific decisions, and accepts full responsibility for the manuscript.

Author Contribution

K.H. was responsible for conceptualization, methodology, investigation, data curation, formal analysis, visualization, writing—original draft, writing—review and editing, and project administration. K.H. approved the final manuscript and accepts full accountability for the work.

Acknowledgements

None.

Data Availability

The anonymized analytical dataset, codebook, source manifest, analysis notebook, and machine-generated outputs required to reproduce the reported analyses are available through the following anonymous OSF view-only link: https://osf.io/zms57/overview?view_only=00393f84839245b39474d27b13c07033. Corporate source documents are not redistributed; public source locators are provided in the source manifest. Internal human-decision records, full-coding registers, and administrative quality-assurance files are excluded from the reviewer repository.

References


1. Abraham, Santhosh. 2007. and Paul Cox. Analysing the Determinants of Narrative Risk Information in UK FTSE 100 Annual Reports. *The British Accounting Review* 39 (3): 227–248. <https://doi.org/10.1016/j.bar.2007.06.002>
2. Adcock, Robert, David Collier. 2001. Measurement Validity: A Shared Standard for Qualitative and Quantitative Research. *American Political Science Review* 95(3):529–546. <https://doi.org/10.1017/S0003055401003100>
3. Beasley, Mark S., and Richard Clune, Dana R. Hermanson. 2005. Enterprise Risk Management: An Empirical Analysis of Factors Associated with the Extent of Implementation. *Journal of Accounting and Public Policy* 24(6):521–531. <https://doi.org/10.1016/j.jaccpubpol.2005.10.001>
4. Bowen, Glenn A. 2009. Document Analysis as a Qualitative Research Method. *Qualitative Research Journal* 9(2):27–40. <https://doi.org/10.3316/QRJ0902027>
5. Boyens, Jon, Angela Smith, Nadya Bartol, Kris Winkler, and Alex Holbrook. 2024. and Matthew Fallon. Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations. NIST Special Publication 800 – 161 Revision 1 Update 1. updated November 1, 2024. <https://doi.org/10.6028/NIST.SP.800-161r1-upd1>
6. Boyson, Sandor. 2014. Cyber Supply Chain Risk Management: Revolutionizing the Strategic Control of Critical IT Systems. *Technovation* 34(7):342–353. <https://doi.org/10.1016/j.technovation.2014.02.001>
7. Campbell, John L., Charles Quincy, Jordan Osseman, K. Ove, and Pedersen. 2013. Coding In-Depth Semistructured Interviews: Problems of Unitization and Intercoder Reliability and Agreement. *Sociological Methods & Research* 42(3):294–320. <https://doi.org/10.1177/0049124113500475>
8. Committee of Sponsoring Organizations of the Treadway Commission. 2017. Enterprise Risk Management—Integrating with Strategy and Performance. COSO. Durham, NC. <https://www.coso.org/enterprise-risk-management>
9. Deumes, Rogier. 2008. Corporate Risk Reporting: A Content Analysis of Narrative Risk Disclosures in Prospectuses. *Journal of Business Communication* 45(2):120–157. <https://doi.org/10.1177/0021943607313992>
10. Doty, D., Harold, and William H. Glick. 1994. Typologies as a Unique Form of Theory Building: Toward Improved Understanding and Modeling. *Academy of Management Review* 19(2):230–251. <https://doi.org/10.2307/258704>
11. European Parliament and Council of the European Union. 2022a. Directive (EU) 2022/2555 on Measures for a High Common Level of Cybersecurity across the Union. *Official Journal of the European Union L* 333:80–152. <https://eur-lex.europa.eu/eli/dir/2022/2555/oj/eng>
12. European Parliament and Council of the European Union. 2022b. Directive (EU) 2022/2557 on the Resilience of Critical Entities. *Official Journal of the European Union L* 333:164–198. <https://eur-lex.europa.eu/eli/dir/2022/2557/oj/eng>
13. Fiss, Peer C. 2007. A Set-Theoretic Approach to Organizational Configurations. *Academy of Management Review* 32(4):1180–1198. <https://doi.org/10.5465/AMR.2007.26586092>
14. Flyvbjerg, Bent. 2006. Five Misunderstandings about Case-Study Research. *Qualitative Inquiry* 12(2):219–245. <https://doi.org/10.1177/1077800405284363>
15. Good, Phillip I. 2005. *Permutation, Parametric, and Bootstrap Tests of Hypotheses*. 3rd ed. New York: Springer. <https://doi.org/10.1007/b138696>
16. Healy, Paul M., and Krishna G. Palepu. 2001. Information Asymmetry, Corporate Disclosure, and the Capital Markets: A Review of the Empirical Disclosure Literature. *Journal of Accounting and Economics* 31(1–3):405–440. [https://doi.org/10.1016/S0165-4101\(01\)00018-0](https://doi.org/10.1016/S0165-4101(01)00018-0)
17. Holm, Sture. 1979. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics* 6(2):65–70. <https://www.jstor.org/stable/4615733>

18. International Organization for Standardization. 2018. ISO 31000:2018 Risk Management—Guidelines. ISO. Geneva.
<https://www.iso.org/standard/65694.html>
19. Jick, Todd D. 1979. Mixing Qualitative and Quantitative Methods: Triangulation in Action. *Administrative Science Quarterly* 24(4):602–611.
<https://doi.org/10.2307/2392366>
20. Krippendorff, Klaus. 2018. *Content Analysis: An Introduction to Its Methodology*. 4th ed. Thousand Oaks, CA: SAGE.
<https://doi.org/10.4135/9781071878781>
21. Linsley, Philip M. 2006. and Philip J. Shrives. Risk Reporting: A Study of Risk Disclosures in the Annual Reports of UK Companies. *The British Accounting Review* 38 (4): 387–404. <https://doi.org/10.1016/j.bar.2006.05.002>
22. Loughran, Tim, Bill McDonald. 2014. Measuring Readability in Financial Disclosures. *Journal of Finance* 69(4):1643–1671.
<https://doi.org/10.1111/jofi.12162>
23. Mikes, Anette. 2009. Risk Management and Calculative Cultures. *Management Accounting Research* 20(1):18–40.
<https://doi.org/10.1016/j.mar.2008.10.005>
24. Mikes, Anette. 2011. From Counting Risk to Making Risk Count: Boundary-Work in Risk Management. *Accounting Organizations and Society* 36(4–5):226–245. <https://doi.org/10.1016/j.aos.2011.03.002>
25. Mikes, Anette, and Robert S. Kaplan. 2015. When One Size Doesn't Fit All: Evolving Directions in the Research and Practice of Enterprise Risk Management. *Journal of Applied Corporate Finance* 27(1):37–40. <https://doi.org/10.1111/jacf.12102>
26. National Institute of Standards and Technology. 2024. The NIST Cybersecurity Framework (CSF) 2.0. NIST Cybersecurity White Paper 29. Gaithersburg, MD. <https://doi.org/10.6028/NIST.CSWP.29>
27. Neuendorf, Kimberly A. 2017. *The Content Analysis Guidebook*. 2nd ed. Thousand Oaks, CA: SAGE. <https://us.sagepub.com/en-us/nam/the-content-analysis-guidebook/book234078>
28. Patton, Michael Quinn. 2015. *Qualitative Research & Evaluation Methods: Integrating Theory and Practice*. 4th ed. Thousand Oaks, CA: SAGE.
<https://us.sagepub.com/en-us/nam/qualitative-research-evaluation-methods/book232962>
29. Posthumus, Shaun, and Rossouw von Solms. 2004. A Framework for the Governance of Information Security. *Computers & Security* 23(8):638–646.
<https://doi.org/10.1016/j.cose.2004.10.006>
30. Schreier, Margrit. 2012. *Qualitative Content Analysis in Practice*. London: SAGE. <https://doi.org/10.4135/9781529682571>
31. Spence, Michael. 1973. Job Market Signaling. *Quarterly Journal of Economics* 87(3):355–374. <https://doi.org/10.2307/1882010>
32. Suchman, Mark C. 1995. Managing Legitimacy: Strategic and Institutional Approaches. *Academy of Management Review* 20(3):571–610.
<https://doi.org/10.2307/258788>
33. Verrecchia, Robert E. 2001. Essays on Disclosure. *Journal of Accounting and Economics* 32(1–3):97–180. [https://doi.org/10.1016/S0165-4101\(01\)00025-8](https://doi.org/10.1016/S0165-4101(01)00025-8)
34. von Solms,, and Basie, Rossouw von Solms. 2004. The 10 Deadly Sins of Information Security Management. *Computers & Security* 23(5):371–376.
<https://doi.org/10.1016/j.cose.2004.05.002>
35. Williams, Paul. 2001. Information Security Governance. Information Security Technical Report 6 (3): 60–70. [https://doi.org/10.1016/S1363-4127\(01\)00309-0](https://doi.org/10.1016/S1363-4127(01)00309-0)

Figures

Figure 1. Firm × domain public-disclosure profiles

Ordered only by sector and neutral firm ID; cell values are human-verified disclosure-specificity codes.




Note: 0 = no qualifying disclosure located in the reviewed corpus; 1 = generic; 2 = named owner/process; 3 = owner/process plus assurance, metric, cadence, or feedback. Codes do not measure practice, performance, compliance, security, or actual maturity.

Figure 1

Firm × domain public-disclosure profile heatmap. Rows are ordered only by observed subgroup and neutral firm_id. Code 0 means not disclosed in the reviewed corpus and does not indicate absence of practice.

Note. Firms are ordered only by observed subgroup and neutral identifier; no firm ranking is encoded.

Figure 2. Sector-level domain profiles in the observed sample



Coverage and uncertainty: Energy: 13 sources, 8/60 flagged cells | Transport: 6 sources, 11/60 flagged cells | Digital infrastructure: 6 sources, 4/60 flagged cells | Banking / FMI: 6 sources, 5/60 flagged cells. Each sector contains six criterion-selected firms. Ranges describe the sample only; they do not estimate sector populations or operational performance.

Figure 2

Observed subgroup domain profiles. Each purposive subgroup contains n=6 firms; medians, ranges, coverage, and uncertainty are descriptive only and do not estimate sector populations.

Note. Firms are ordered only by observed subgroup and neutral identifier; no firm ranking is encoded.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [TableA12.docx](#)