

**A Multi-Source Statistical Framework
for Adequate Organizational Arrangements
Assessment under the Italian Crisis Code**

Giuseppe Marino *, Alessandra Mancini *, Gianluca Egidi **,
Andrea Colantoni ****, Massimiliano Meschini ** and Matteo Campora *****

* Risk management consultant - Assytech s.a.s., via dodici ottobre 28 r - 16121 - Genova (Italy)

** SEA Tuscia S.r.l., via San Camillo De Lellis snc, 01100 Viterbo (Italy) and Aleph S.r.l., via San Carlo, 01100 Viterbo (Italy)

*** Avvocato, Foro di Genova (Italy)

**** Aleph S.r.l, via San Carlo 32, 01100 Viterbo (Italy)

Corresponding author: Giuseppe Marino

Contents

1	Risk Check: A Multi-Source Statistical Framework for Adequate Organizational Arrangements Assessment under Italian Crisis Code (CCII) and D.Lgs. 47/2026	2
1.1	Abstract.....	3
1.2	1. Introduction.....	3
1.3	2. Related Work and Regulatory Context	4
1.3.1	2.1 Italian regulatory framework.....	4
1.3.2	2.2 Commercial state of the art	5
1.3.3	2.3 Academic precedents.....	6
1.4	3. System Architecture	6
1.4.1	3.1 Identity, persistence and infrastructure	7
1.4.2	3.2 Reverse proxy and CI/CD	8
1.5	4. Methodology	8
1.5.1	4.1 The Exposure Index (IEA)	8
1.5.2	4.2 Context-aware Kc calibration	9
1.5.3	4.3 Statistical harmonization $D_{14} \times D_{61}$	9
1.5.4	4.4 Information-gain question selection.....	10
1.5.5	4.5 Procedural D14 generation from enterprise profile.....	10
1.5.6	4.6 LLM-augmented Nota Integrativa analysis.....	11
1.5.7	4.7 Real DSCR computation under CCII art. 13	11
1.5.8	4.8 Adaptive question ordering — runtime financial-signal driven	12
1.5.9	4.9 CERV four-dimensional response vector — D61 full battery	13
1.5.10	4.10 Self-reporting bias detection.....	14
1.5.11	4.11 Internal consistency index	14
1.5.12	4.12 Bayesian posterior update of the IEA via documentary evidences.....	15
1.5.13	4.13 Sensitivity analysis — per-question IEA leverage	16
1.5.14	4.14 Greedy remediation optimiser — the A.S.S.E.T. plan	17
1.5.15	4.15 Inter-question correlation matrix — independence assumption validation.....	17
1.6	5. Implementation	18
1.6.1	5.1 Technology stack	18
1.6.2	5.2 Self-hosted infrastructure	18
1.6.3	5.3 Multi-tenancy and authentication.....	19
1.6.4	5.4 D.Lgs. 47/2026 information flows.....	19
1.6.5	5.5 AI coaching companion — contextual feedback during the interview.....	19
1.7	6. Statistical Validation.....	20
1.8	7. Discussion.....	21
1.8.1	7.1 The intentional omission of sectoral benchmarking.....	21
1.8.2	7.2 The role of professional judgement	21
1.8.3	7.3 Comparison with the dominant commercial product.....	21
1.9	8. Future Work	22
1.10	9. Conclusion	23
1.11	Acknowledgments.....	24
1.12	References	24

1.1 Abstract

The Italian Crisis and Insolvency Code (Codice della Crisi d'Impresa e dell'Insolvenza, *CCII*, art. 2086 c.c.) and the more recent D.Lgs. 47/2026 (“Legge Capitali”) impose on the directors of every legal entity the obligation to establish and continuously verify *adequate organizational, administrative and accounting arrangements* for the early detection of corporate distress. While the obligation is enforced by jurisprudence (Cassazione 4849/2023) and by case-law in commercial sections nationwide, no software platform in the Italian market currently produces an evidence-grade, statistically validated assessment that combines (i) the structured financial data deposited in the Registro Imprese (XBRL), (ii) the narrative content of the explanatory notes (Nota Integrativa), and (iii) the qualitative answers collected through expert interviews. We present **Risk Check**, an end-to-end web application that fuses these three information sources and produces a probative-quality assessment, traceable to art. 2381-ter, 2403 and 2409-bis under D.Lgs. 47/2026. The platform introduces nine methodological contributions: (1) a statistical predictor that harmonizes a short screening questionnaire (14 items, *D14*) with a comprehensive instrument (63 items, *D61*), achieving root-mean-squared error (RMSE) below 0.05 and Pearson correlation above 0.80 on a Monte Carlo benchmark of 1,000 synthetic respondents; (2) a context-aware Hidden Gap Coefficient (*Kc*) calibration grounded on five Italian regulatory contexts (general art. 2086 *CCII*, Gelli-Bianco healthcare, food traceability, logistics, third sector); (3) an inference-based dynamic question generator that derives a bespoke *D14* battery from the 63-item *CNDCEC* corpus by combining boost rules conditioned on the enterprise profile (sector, *ATECO*, legal form, headcount, revenue trend, credit rating) with a runtime adaptive ordering driven by live financial-deterioration signals; (4) a large-language-model-augmented analysis of the explanatory notes that produces a structured, citable JSON; (5) a four-dimensional *CERV vector* assessment framework for the full *D61* battery that decomposes each response into Compliance, Evidence, Risk, and Uncertainty coordinates with hard-blocking rules for critical governance gaps; (6) a suite of answer-quality diagnostics — a Self-Reporting Bias Index grounded on Paulhus (1984) socially desirable responding theory, an internal-consistency coefficient inspired by Cronbach (1951) alpha, a Bayesian Beta-Binomial posterior update of the IEA calibrated on uploaded documentary evidences, and a one-at-a-time sensitivity analysis (Tornado chart); (7) a greedy *Quick Win / Piano A.S.S.E.T.* optimizer that identifies the remediation sequence with maximum efficiency (Δ IEA per consultant-day), with a seven-step IEA waterfall chart and alignment to ISO 31000:2018 §6.5.3; (8) a 14×14 Pearson correlation matrix generated via Monte Carlo simulation (N=3,000, deterministic seed 42), with a two-level latent structure and Union-Find clustering for highly-correlated pairs, which validates the IEA independence assumption and identifies informationally redundant questions; (9) a real-time AI coaching companion embedded in the questionnaire wizard that provides normatively grounded commentary per question category with a traceable coaching log useful for Business Judgement Rule defence. All components are released as a fully self-hosted stack (Hetzner Cloud, no Big Tech dependency) compatible with the EU data-sovereignty constraints faced by Italian SMEs and professional service firms.

1.2 1 Introduction

The Italian reform of company law in 2019 (D.Lgs. 14/2019), followed by the more recent D.Lgs. 47/2026 (commonly referred to as “Legge Capitali”), introduces a paradigm shift in directorial liability. The obligation to “establish an organizational, administrative and accounting setup adequate to the nature and size of the enterprise” (art. 2086 c.c.) has moved from a static, document-based compliance to a *continuous behavioural obligation*: directors must not only have the documents — they must demonstrate that information flows from operational layers up to the board of directors, the supervisory body and the statutory auditor, and that those organs receive, discuss and minute the information they receive (D.Lgs. 47/2026, art. 2381-ter; ICAS, art. 13).

The market response to this regulatory shift has been uneven. A small number of commercial platforms (here anonymised; see Appendix D) automate the extraction of structured financial data from the XBRL deposits and

produce visual dashboards (Altman Z-Score, K-Score, prospective DSCR). However, an independent review of these products (this work, §3) reveals two systematic shortcomings: (a) the narrative explanatory notes — which under Italian accounting law contain the only disclosure of debt-installment schedules required to compute the *real* Debt Service Coverage Ratio under CCII art. 13 — are not read at all; (b) no platform produces an assessment that maps to the three pillars of art. 2086 (organizational / administrative / accounting arrangements) and that could be admitted in evidence in a corporate liability proceeding.

This paper describes Risk Check, a software platform developed by the authors. to fill this gap. The platform combines structured XBRL parsing, large-language-model-augmented analysis of the explanatory notes, two qualitative diagnostic batteries (a 14-item short form *D14* and a 63-item long form *D61* derived from the CNDCEC professional checklist), and a domain-specific health-sector instrument (the *Gelli-Bianco* battery, Law 24/2017). The output is a 10-section A4 report, digitally sealed with a TSA-compliant timestamp (Aruba PEC), and a parallel log of inter-organ information flows whose cryptographic hash (SHA-256) provides probative immutability under art. 2381-ter D.Lgs. 47/2026.

The paper's contributions are organized as follows: §2 surveys the regulatory and commercial context. §3 introduces the four-layer architecture. §4 details the statistical methodology across fifteen sub-sections: the Exposure Index and its Kc calibration (§4.1–4.2), the *D14*–*D61* harmonisation predictor (§4.3), the Information Gain selection framework (§4.4), the contextual *D14* generator (§4.5), the adaptive question ordering driven by real-time financial signals (§4.8), the CERV four-dimensional answer-quality vector (§4.9), the Self-Reporting Bias Index (§4.10), the internal consistency measure (§4.11), the Bayesian Beta-Binomial posterior update via documentary evidences (§4.12), the per-question sensitivity analysis (§4.13), the Quick Win / Piano A.S.S.E.T. optimizer (§4.14), the 14× 14 Pearson correlation matrix (§4.15), the LLM-augmented NI reader (§4.6), and the real DSCR computation (§4.7). §5 describes the implementation choices that enable the platform to operate as a fully self-hosted, EU-resident service. §6 reports the Monte Carlo validation. §7 discusses limitations and §8 outlines future work.

1.3 2. Related Work and Regulatory Context

1.3.1 2.1 Italian regulatory framework

The duties of directors with respect to organizational adequacy are codified in art. 2086 c.c., paragraph 2:

The entrepreneur operating in a corporate or collective form has the duty to establish an organizational, administrative and accounting setup that is adequate to the nature and size of the enterprise, also in the perspective of the timely detection of the crisis of the enterprise and of the loss of going-concern.

The Italian Crisis and Insolvency Code (ICAS, D.Lgs. 14/2019, in force since 15 July 2022) further specifies five early-warning indicators (art. 13): sustainability of debt-installments over the next six months (the *real* Debt Service Coverage Ratio), debt-to-equity adequacy, repeated tax irregularities, social-security contribution arrears, and trade-payable arrears beyond ninety days. Failure to react to a deterioration of any of these indicators may give rise to directors' personal liability under art. 2476 c.c. (Cass. civ. 4849/2023).

D.Lgs. 47/2026 (“Legge Capitali”, entered into force April 2026) introduces three new mandatory information flows:

Flow	Source → Destination	Frequency	Legal basis
1	Delegated bodies → Board of Directors	quarterly (monthly if DSCR < 1)	art. 2381-ter (new)
2	Board of Directors → Statutory Auditor	quarterly + annual	art. 2403 (strengthened)

Flow	Source → Destination	Frequency	Legal basis
3	System → Legal Auditor	annual	art. 2409-bis (extended)

The compliance obligation is no longer documental but behavioural: directors must demonstrate that information has been **transmitted, received and discussed** — not merely produced.

1.3.2 2.2 Commercial state of the art

The leading commercial product in the Italian adequate-arrangements market — commercially available since 2018, anonymised here as **Vendor A** for the reasons discussed in Appendix D — produces an automated report from the *partita IVA* (VAT number) alone. Its pipeline is: VAT number → financial data API (Cerved or AIDA) → XBRL parsing → ratio computation → PDF template. The output features Z-Score, K-Score and a prospective DSCR.

This work performs an independent feature-level comparison (Table 1):

Function	Vendor A	Risk Check
XBRL ratio computation	✓	✓
Four-year time series	✓	✓ (subject to data provider contract)
Altman Z-Score / K-Score	✓	✓
Sectoral benchmark	✓	deliberately omitted (see §7)
Reading of Nota Integrativa	×	✓ (LLM-augmented)
Real DSCR (ex CCI art. 13)	not computable	computed from NI text
Mapping to art. 2086 (three pillars)	×	✓
D.Lgs. 47/2026 information flows	×	✓ (with SHA-256 immutability)
Evidentiary value (BJR)	×	✓ (TSA-sealed report)
Context-adaptive question battery	×	✓ (sector × ATECO × credit signal)
Runtime question priority ordering	×	✓ (live financial-deterioration signals)
CERV four-dimensional answer quality	×	✓
Self-reporting bias detection		(E-coverage, A-bias, U-mean composite)
Internal consistency index	×	✓ (Cronbach-inspired, per ISO 31000 area)
Bayesian IEA posterior (evidences)	×	✓ (Beta-Binomial, CI 95%, trace chart)
Sensitivity / Tornado chart	×	✓ (per-question IEA leverage, Pareto index)
Quick Win optimizer (Piano A.S.S.E.T.)	×	✓ (greedy Δ IEA/consultant-days, 7-step waterfall)

Function	Vendor A	Risk Check
14x 14 question correlation matrix	×	✓ (Monte Carlo N=3,000, Union-Find clustering)
In-wizard AI coaching companion	×	✓ (6 categories, 10-min cache, BJR log)

The two distinguishing properties of Risk Check — the reading of the explanatory notes and the production of an evidence-grade dossier — derive from a deliberate methodological choice that places professional judgement (the *terzo controllore*) at the centre of the system, rather than substituting it with an algorithm. Sectoral benchmarking is omitted because the comparison with a synthetic peer-group is, in our view, normatively irrelevant: art. 2086 imposes adequacy with respect to the *nature and size of the enterprise itself*, not with respect to a population average.

1.3.3 2.3 Academic precedents

The use of LLMs to extract structured information from financial disclosures is a recent and rapidly growing field. Sun et al. (2024) demonstrated zero-shot extraction of risk factors from 10-K filings using GPT-4; Wang & Lou (2025) reported a 92% F1 score on a similar task with Claude Sonnet 3.5. The application to Italian XBRL deposits — encoded under the Italian *tassonomia itcc-ci-2018-11-04* rather than US-GAAP — has, to our knowledge, not been previously documented.

Information-theoretic question selection in psychometrics has a long tradition (Shannon, 1948; Quinlan, 1986; Hambleton et al., 1991). We adapt it here to a regulatory-compliance setting where the latent variable is an exposure index rather than a personality or proficiency trait. The use of Bayesian regression with bootstrap confidence intervals (Efron & Tibshirani, 1986) to harmonise instruments of different length is, again to our knowledge, novel in the corporate-compliance literature.

The application of *computerized adaptive testing* (CAT) principles — where the item selection is conditioned on previously observed responses and on examinee profile characteristics — has an extensive literature in educational measurement (Wainer, 1990; van der Linden & Glas, 2000). Risk Check’s adaptive-ordering module (§4.8) draws on this tradition but adapts it to a regulatory interview context: rather than maximising information about a latent ability trait, it maximises legal-exposure coverage per unit of interview time, using the enterprise’s sector, ATECO classification, and financial-deterioration signals as conditioning variables.

The detection of self-reporting bias in self-administered questionnaires is grounded in the two-component model of socially desirable responding (Paulhus, 1984): *self-deceptive enhancement* (respondents genuinely believe overstated claims) and *impression management* (respondents consciously overstate). Our bias index (§4.10) focuses on observable proxies — evidence coverage, declared-adequacy rate, uncertainty mean — rather than latent psychological constructs, making it operationalisable without dedicated psychological measurement instruments. The FNC 25/07/2023 professional reference document provides an empirical baseline of expected compliance rates for Italian SMEs that is used as the benchmark for the “atypical adequacy-rate” warning. The Bayesian Beta-Binomial posterior update of the IEA (§4.12) follows the standard conjugate-prior framework (Box & Tiao, 1973) and is novel in the regulatory-technology domain to our knowledge.

1.4 3. System Architecture

The Risk Check platform is organized in four layers (Fig. 1). All components are deployed on a single Hetzner Cloud dedicated server in Germany (EU data residency).

Layer	Components
Layer 4 — Distribution	PDF Assessment Report (10 sections, A4, TSA-sealed) · D.Lgs. 47/2026 information flows (AD→CDA, CDA→Sindaci, Sistema→Revisore) · Email transport with SHA-256 immutability log
Layer 3 — Analysis	Exposure Index IEA (Kc, 5 regulatory contexts) · D14× D61 statistical harmonisation predictor · Contextual D14 generator (boost rules × ATECO/sector) · Adaptive question ordering (runtime financial signals) · Information Gain selection (Shannon entropy) · CERV vector [C,E,R,U] (VITAL/IMPORTANT/ACCESSORY clusters) · Self-Reporting Bias Index · Internal Consistency Index (Cronbach-inspired, per area) · Bayesian IEA update (Beta-Binomial posterior per evidence) · Sensitivity Analysis / Tornado chart · Quick Win optimizer / Piano A.S.S.E.T. (greedy Δ IEA/consultant-days) · Correlation Matrix 14×14 (Monte Carlo N=3,000) · CCII Indicators (current ratio, quick ratio, DSCR proxy & real, PFN/EBITDA, ICR) · LLM-augmented Nota Integrativa Reader (Claude Sonnet 4.6)
Layer 2 — Data Ingestion	XBRL parser (tassonomia itcc-ci-2018-11-04) · Three diagnostic batteries (D14, D61, Gelli-Bianco) · AI coaching companion (6 categories, 10-min cache, BJR log) · Evidence upload (S3-compatible, signed URLs, anti-CSRF)
Layer 1 — Identity & Persistence	Auris IAM (Keycloak 26.4, OIDC, RS256, multi-tenant) · PostgreSQL 16 (Prisma ORM, 19 entities) · MinIO object storage (S3-compatible) · Caddy reverse proxy (DNS-01 ACME, hardened nftables)

Fig. 1: The four-layer architecture of Risk Check.

1.4.1 3.1 Identity, persistence and infrastructure

Authentication is delegated to *Auris*, an enterprise-grade Keycloak-based identity platform co-deployed on the same infrastructure. *Auris* implements OpenID Connect (OIDC) with RS256 token signing, realm-level isolation per tenant, brute-force lockout, multi-factor authentication, and audit logging of all authentication events — a security posture consistent with OWASP ASVS Level 2 for IAM and with the accountability requirements of GDPR art. 5(1)(f). OpenID-Connect tokens (RS256) are validated against a JWKS endpoint cached locally; the multi-tenancy boundary is enforced at the SQL level by a mandatory `tenant` column on every entity, and at the storage level by service accounts scoped to per-tenant object storage buckets (`assy-evidenze`, `assy-documenti-identita`, `assy-reports-sealed`).

The database is a single PostgreSQL 16 instance with separate logical databases per tenant. The schema is managed via Prisma 5.22 with version-controlled migrations applied during deployment by an isolated migrator container. The choice of a fully self-hosted infrastructure on EU territory (Hetzner DE, Falkenstein) was driven by two considerations: (i) the data processed by the platform — financial statements, qualitative interviews, evidence documents — is intrinsically sensitive under GDPR; (ii) the obligation to provide an audit trail of who saw what, and when, is incompatible with the indirect access models of public-cloud Big Tech providers.

1.4.2 3.2 Reverse proxy and CI/CD

A Caddy reverse proxy handles TLS termination via the DNS-01 ACME challenge, enabling automatic certificate management for the *.assy.tech domain space without requiring any inbound HTTP plaintext port. The application is served exclusively over HTTPS, with TLS 1.2/1.3 enforced at the proxy layer and HSTS pre-loading activated. The Next.js application container is not exposed directly to the Internet.

Continuous deployment is performed by a self-hosted CI runner co-located on the same machine. The deployment script is idempotent and atomic: it pulls the source code, builds the container image, applies pending schema migrations, recreates the service container, and performs a health-check probe against the production login endpoint. Failure at any step aborts the deployment without rolling forward.

Cron jobs are implemented as systemd timers on the host, not as application-level schedulers, in order to decouple their lifecycle from the application container's. Two cron services are currently active: `assy-scadenza-report.timer` (daily at 09:00 UTC, sends an alert seventy-five days before the report expiration) and `assy-notifiche-adeguamento.timer` (daily at 07:00 UTC, reminds operators of pending evidence uploads).

1.5 4. Methodology

This section formalises the four methodological contributions of Risk Check: (i) the *Exposure Index* (IEA) with its Hidden Gap Coefficient K_c ; (ii) the statistical harmonisation between the 14-item screening battery and the 63-item full battery; (iii) the information-gain-based selection of cardinal questions; (iv) the LLM-augmented analysis of the explanatory notes.

1.5.1 4.1 The Exposure Index (IEA)

Let Q be the set of binary-trichotomous diagnostic questions (*SI, PARZIALE, NO*, encoded as 2, 1, 0). For a respondent providing answer $a_i \in \{0, 1, 2\}$ to question i , define the *light gap* per question as a monotonically-decreasing affine map:

$$g_i = \text{answerToGap}(a_i) = \begin{cases} 0 & a_i = 2 \\ 0.5 & a_i = 1 \\ 1 & a_i = 0 \\ 0.75 & \text{(unknown / cautionary default)} \end{cases}$$

Questions are organised in five domains of the ISO 31000 framework, labelled *A* through *E*. The *area gap* is the arithmetic mean of the per-question gaps within an area; if an area has no responses, the gap is set to 0 (no signal).

The *light gap* is then calibrated with a *Hidden Gap Coefficient* (K_c), which estimates the residual risk not captured by the limited number of sentinel questions:

$$\hat{g}_a = g_a + (1 - g_a) \cdot K_c(a)$$

The K_c coefficients (Table 2) are estimated separately for each regulatory context and for each area. The values reflect the conservative principle (*principio di cautela*) of Italian commercial jurisprudence: when in doubt, assume more risk, not less.

Table 2 — Hidden Gap Coefficients (general context, art. 2086 CCII baseline):

Area	Description	Kc	Confidence
A	Organizational arrangements	0.18	0.82
B	Administrative arrangements	0.28	0.72
C	Accounting arrangements	0.25	0.75
D	Risk management & continuity	0.35	0.65
E	Insurance & risk transfer	0.15	0.85

The Exposure Index is finally computed as a weighted average over the five areas, with weights w reflecting the legal exposure of the director:

$$IEA = \frac{\sum_{a \in \{A..E\}} \hat{g}_a \cdot w_a}{\sum_a w_a}$$

The weights are $w_A = w_B = w_C = 2, w_D = 3, w_E = 1$: area D (risk management and going concern) is given triple weight because its failure most directly triggers art. 2476 c.c. liability. The resulting $IEA \in [0,1]$ is classified on a four-level qualitative scale (Table 3):

Table 3 — IEA qualitative classification:

IEA range	Level
[0.00, 0.25]	LOW
(0.25, 0.55]	MEDIUM
(0.55, 0.80]	SIGNIFICANT
(0.80, 1.00]	HIGH

1.5.2 4.2 Context-aware Kc calibration

The same five-area framework is applied to four additional regulatory contexts (Gelli-Bianco healthcare under Law 24/2017, food traceability under Reg. UE 178/2002, logistics under CMR convention, third sector under D.Lgs. 117/2017). Each context redefines the Kc coefficients to reflect its specific risk landscape. For example, the Gelli-Bianco context increases Kc(A) (organizational adequacy) to 0.28 because hospital governance under Italian law requires a *Risk Manager* with formal appointment — an absence which the general context would penalise less severely.

The weight vector w is also context-dependent. The Manuale Risk Check v1.0 (internal document, May 2026) provides a five-by-five matrix (five sections× five contexts) that is the single source of truth for the platform’s calibration.

1.5.3 4.3 Statistical harmonization D14 × D61

The 14-item battery *D14* is conceived as a *cardinal subset* of the 63-item *D61* battery (derived from the CNDCEC professional checklist). For each *D14* question q , a mapping $M(q)c_{D61}$ identifies the one to six *D61* items it is meant to subsume. The mapping was manually constructed by reading both batteries; it is exposed in the source code as a constant and is amenable to expert review.

Given a *D14* answer vector, the platform computes two quantities: the *D14* exposure index (using the formula of §4.1) and a *predicted D61 exposure index*. The prediction algorithm proceeds in three steps.

Step 1 — Per-question contribution. For each *D14* question q , the *fan-out* $|M(q)|$ — the number of *D61* questions it covers — determines a contribution coefficient

$$\beta_q = 1 + \log_2 |M(q)|$$

reflecting the fact that a *D14* question with high fan-out is a more reliable proxy for the underlying *D61* signal than a question with low fan-out.

Step 2 — Area aggregation. Per-question contributions are aggregated by area (using the same A–E mapping of §4.1), weighted by the calibrated gap:

$$g_a^{D61} = \frac{\sum_{q \in a} \text{calibrateGap}(g_q, a) \cdot \beta_q}{\sum_{q \in a} \beta_q}$$

Step 3 — IEA composition. The predicted *D61* index is then composed using exactly the same area weights w of §4.1:

$$\tilde{\text{IEA}}^{D61} = \frac{\sum_a g_a^{D61} \cdot w_a}{\sum_a w_a}$$

This *structural isomorphism* between the *D14*-derived prediction and the *D61* baseline is critical: it guarantees that, when the predictor is benchmarked against the *D61* formula applied to a hypothetical full response, the systematic bias is bounded by the residual of the per-question calibration alone.

Confidence intervals. A bootstrap-style 95% confidence interval is reported alongside the point estimate, with the standard error derived from a measure of *predictor confidence*:

$$c = \frac{1}{|D14|} \sum_q \min\left(1, \frac{|M(q)|}{4}\right)$$

and the standard error empirically scaled as $\sigma = 0.15 \cdot (1 - c)$. Future versions will substitute the empirical scaling with a true bootstrap resampling once a corpus of paired *D14/D61* responses becomes available.

1.5.4 4.4 Information-gain question selection

Following Shannon (1948) and Quinlan (1986), we compute the *information gain* of each *D14* question with respect to the IEA as discrete latent variable:

$$\text{IG}(Q_i) = H(\text{IEA}) - H(\text{IEA} | Q_i)$$

where $H(X)$ is the Shannon entropy of X discretised in the four bins of Table 3 and $H(X | Q_i)$ is the conditional entropy of X given Q_i . The computation is performed by Monte Carlo over the $3^{14} = 4,782,969$ possible response patterns of *D14*, sampled with a deterministic seed (Mulberry32). The top k questions by IG are exposed in the summary view as “cardinal questions” and used to prioritise follow-up evidence collection.

1.5.5 4.5 Procedural D14 generation from enterprise profile

A practical objection to a one-size-fits-all *D14* is that the cardinal questions for a logistics SME (e.g., CMR liability, cold-chain risk) differ from those for a healthcare provider (clinical risk, ERM under Gelli-Bianco). Risk Check implements a contextual *D14 generator* that, given an enterprise profile (sector, revenue band, legal form, ATECO code, headcount, recent revenue variation, credit rating), reweights the *D61* questions according to a set of expert-curated boost rules:

- *Healthcare sector*: × 1.5 boost on questions 27, 28, 30, 31, 39, 62, 63 (SCIGR, OdV 231, whistleblowing, clinical risk).
- *Manufacturing / Food*: × 1.4 boost on questions 33, 39, 47, 56 (certifications, internal audit, integrated accounting).
- *Capital companies (S.r.l. / S.p.A.)*: × 1.3 boost on questions 36–42 (board of directors, internal audit, industrial plan).
- *Revenue decline > 5%*: ×1.6 boost on questions 54, 55, 58, 59, 60, 61 (crisis indicators, KPI, treasury).
- *Critical credit rating (D/E/F)*: ×1.8 boost on questions 54, 55, 59, 61.

The top 14 *D61* questions, after the boost, are selected as the *contextual D14*. A 20% additional bonus is applied to questions that belong to the baseline *D14* cardinal subset, in order to preserve a coverage floor. The generated battery is presented to the operator with a per-question justification (“Healthcare sector requires ERM + 231 + clinical risk”), which is in turn reproduced in the final assessment as evidence of the bespoke approach.

1.5.6 4.6 LLM-augmented Nota Integrativa analysis

The explanatory notes of an Italian financial statement are embedded in the XBRL deposit as escaped HTML within `<itcc-ci:Commento*>` elements. The platform extracts these elements, decodes the entities, strips the HTML markup, and concatenates the textual content into a single document of typically 5,000–25,000 words.

The document, together with the structured CE/SP, is sent to the Claude Sonnet 4.6 API with a system prompt that requires the model to return *exclusively* a JSON object conforming to the *NIAAnalysis* interface (Listing 1):

```

1 interface NIAAnalysis {
2   going_concern: { dichiarato: boolean; testo: string; incertezze: boolean }
3   clausola_2086: { presente: boolean; testo: string; gap: string[] }
4   coerenza: { coerente: boolean; problemi: string[] }
5   dscr_dati: { debiti_correnti_finanziari: number | null; quote_capitale_6m: number |
6     null }
7   compensi_organici: { ammontare: number | null; anomalia: boolean }
8   partecipazioni: Array<{ nome: string; valore: number | null; rischio: boolean }>
9   rischi_nascosti: string[]
10  gap_disclosure_47_2026: string[]
11  rating_ni: 'ADEGUATA' | 'PARZIALE' | 'INADEGUATA' | 'NON_DISPONIBILE'
12  tokens_usati?: number

```

Three architectural choices distinguish this implementation from a naive prompt-and-pray approach. First, the system prompt is *domain-specialised*: it instructs the model to apply Italian corporate law (art. 2086 c.c., D.Lgs. 47/2026) and to cite verbatim the passages it relies on, which enables a downstream legal audit of the model’s reasoning. Second, the output is *strictly typed*: any JSON parse failure triggers a re-execution with a stricter system prompt; any missing field is replaced with a conservative default (`NON_DISPONIBILE` rating, empty gap lists). Third, the analysis is *not* a substitute for professional judgement: the JSON is rendered in the assessment as one section among ten, and the *terzo controllore* (the consultant operating the platform) signs off on it.

1.5.7 4.7 Real DSCR computation under CCII art. 13

The Debt Service Coverage Ratio mandated by CCII art. 13 is defined as the ratio between operating cash flow and the principal-plus-interest instalments due within the next six months. The first quantity is approximated by EBITDA (sum of operating margin and depreciation, as parsed from the XBRL CE); the second quantity is *not* available in the structured XBRL elements — it is disclosed only in the textual portion of the explanatory notes (“scadenze dei debiti finanziari”).

Risk Check is, to our knowledge, the first Italian platform that computes the *real* DSCR by extracting `quote_capitale_6m` from the explanatory notes via the LLM reader of §4.6, and combining it with the parsed `oneri_finanziari` and `ebitda`:

$$DSCR_{\text{real}} = \frac{\text{EBITDA}}{2 \cdot \text{quote_capitale_6m} + \text{oneri_finanziari}}$$

When the LLM-derived `quote_capitale_6m` is null (notes do not disclose, or disclose ambiguously), the platform falls back to a *proxy DSCR*:

$$DSCR_{\text{proxy}} = \frac{\text{EBITDA}}{0.30 \cdot \text{debiti_entro_esercizio} + \text{oneri_finanzia}}$$

where the 0.30 multiplier estimates the financial fraction of short-term debt. Both quantities are reported with a traffic-light classification (≥ 1.5 green, ≥ 1.0 yellow, otherwise red) and with a textual rationale that the consultant can edit before report finalisation.

1.5.8 4.8 Adaptive question ordering — runtime financial-signal driven

The contextual *D14* generator of §4.5 operates *before* the interview session: it derives the bespoke battery from the 63-item CNDCEC corpus using static profile data (sector, legal form, ATECO code, headcount, declared revenue band). A complementary, distinct mechanism operates *during* the session: the `adaptive-selection` module computes a priority score for each question at render time, incorporating live financial signals obtained from the enterprise’s most recent XBRL balance sheets — signals that are not available when the questionnaire is first composed.

The priority score for question q belonging to area a is defined as:

$$\text{priority}(q) = \pi_a \cdot w_q \cdot \delta_{\text{sentinel}}(q)$$

where $\pi_a \in [0, 1]$ is the area priority computed from the financial signal profile, w_q is the base weight of the question (from the mapping JSON), and $\delta_{\text{sentinel}}(q) = 1.5$ if q is designated a sentinel question, 1 otherwise.

Area priorities π_a are computed via an additive boost model applied to a uniform baseline of 0.5:

Financial signal	Affected area(s)	Boost magnitude
Revenue decline $\geq 5\%$ YoY	D (risk management & going-concern)	+0.25
Revenue growth $> 10\%$ + profit decline	D	+0.30
Credit rating D or E	D	saturation at 1.0
Credit score < 50	D	+0.30
ATECO 86–88 (healthcare)	A, E	+0.15 / +0.20
ATECO 49–53 (transport)	D, E	+0.15 each
ATECO 62–63 (ICT)	E	+0.20
Sector: manufacturing / food	B, C	+0.25 / +0.20
Sector: technology / ICT	E, B	+0.30 / +0.15
Sector: third sector / cooperative	A, B	+0.25 / +0.20

Boosts from sector-text and ATECO-code rules compose multiplicatively when both fire: an ATECO 86 entity in the “sanitary” sector accumulates both the text-match and the ATECO boost. The design ensures that the area

weights fed into the IEA formula (§4.1) and the question ordering seen by the operator are *coherent*: areas under financial stress receive both higher Kc penalty and earlier attention during the interview.

The output of *adaptive-selection* is used in two ways. First, the questionnaire wizard presents questions in order of descending priority, so that an operator conducting a time-limited interview always covers the most legally-exposed areas first. Second, the top five questions are flagged with a “Priority for your sector” badge, drawing the operator’s attention to the questions whose answers have the highest expected impact on the final assessment.

This *runtime adaptivity* is architecturally distinct from the *pre-session procedural generation* of §4.5: the two pipelines are composable — the contextual D14 selects *which* D61 items to include, while the adaptive sorter determines *in which order* to present them. In the default configuration the 14-item battery is first generated from D61 (§4.5) and then sorted by the adaptive module before display.

1.5.9 4.9 CERV four-dimensional response vector — D61 full battery

For the 63-item *D61* battery (conducted as a deeper follow-up interview), each answer is encoded as a four-dimensional vector rather than a scalar gap:

$$\mathbf{v}_i = (C_i, E_i, R_i, U_i)$$

where the four dimensions are defined as follows.

C — Compliance is the degree of literal adherence to the statutory requirement: $C_i = 1.0$ when the respondent answers *SI* and declares adequacy; $C_i = 0.5$ when *SI* but declares “non adeguato” (present but insufficient); $C_i = 0.0$ for *NO*; and $C_i = 1.0$ for *N/A* (legally inapplicable).

E — Evidence is the degree to which the declared compliance is supported by documentary evidence: $E_i = 1.0$ when a reference document is cited; $E_i = 0.3$ when the answer is *SI* but no document is cited (self-assertion without proof); $E_i = 0.0$ for *NO*.

R — Risk is the weighted impact of non-compliance on going-concern continuity:

$$R_i = (1 - C_i) \cdot w_{\text{cluster}(i)}$$

where w_{cluster} reflects the triclass criticality scheme (*vital* = 0.7, *important* = 0.2, *accessory* = 0.1), derived from the FNC 25/07/2023 professional document and the CNDCEC 2025 conduct standards.

U — Uncertainty captures the epistemic ambiguity of the answer: $U_i = 0.0$ for *N/A*; $U_i = 0.1$ for a clean *NO*; $U_i = 0.3$ for *SI-adequate* without documentary evidence; $U_i = 0.3-0.6$ for *SI-non-adequate* (higher when no document is cited); $U_i = 0.0$ for *SI-adequate* with evidence.

The aggregate cluster score is computed as:

$$\text{score}_{\text{cluster}} = \frac{1}{n_{\text{E cluster}}} \sum_{\text{E cluster}} C_i \cdot (0.7 + 0.3 \cdot E_i)$$

This formulation ensures that *compliance without evidence* ($E = 0.3$) contributes only 79% as much as *compliance with evidence* ($E = 1.0$) to the cluster score — a deliberate penalty that reflects the evidentiary standard of Italian commercial proceedings.

The composite score is a convex combination of the three cluster scores:

$$\text{score}_{\text{composito}} = 0.7 \cdot \text{score}_{\text{vital}} + 0.2 \cdot \text{score}_{\text{important}} + 0.1 \cdot \text{score}_{\text{accessory}}$$

Hard-blocking rules apply before the composite score is used: (a) any *vital* question answered *NO* unconditionally produces a verdict of *INADEGUATO*, regardless of the aggregate; (b) any *important* question answered *NO* produces *ADEGUATO CON RISERVE*; (c) any *vital* question with $U \geq 0.4$ triggers an automatic request for independent data integration. These blocking rules encode the principle that certain governance gaps — absence of a board of directors, absence of a legal auditor in a *S.r.l.* above the statutory threshold, absence of treasury monitoring — are *categorical* failures under Italian law and cannot be compensated by excellence in other areas.

The CERV framework thus extends the IEA scalar of §4.1 into a multi-dimensional diagnostic: the IEA tells *how much* exposure exists; the CERV matrix tells *where* the exposure is rooted and *how much uncertainty* surrounds the measurement.

1.5.10 4.10 Self-reporting bias detection

A well-documented limitation of self-administered compliance questionnaires is the *self-serving bias* (also termed *social desirability responding*): respondents tend to over-claim compliance, over-rate adequacy, and under-declare uncertainty (Paulhus, 1984). In a regulatory-technology context this is not merely a measurement nuisance — it is a *probative risk*: a director who relies on an assessment based on systematically over-stated answers may face aggravated personal liability.

Risk Check addresses this via an explicit *Self-Reporting Bias Index* computed from three observable proxies:

1. **E-coverage** (ε) : the fraction of *SI* answers supported by a cited documentary evidence. An honest and well-prepared respondent is expected to be able to document the majority of their *SI* answers. The critical threshold is $\varepsilon < 0.50$.
2. **Declared-adequacy rate** (α_d) : the fraction of *SI* answers flagged as *adeguato* (fully adequate). The FNC 25/07/2023 document reports an expected range of 50–70% for Italian SMEs; a rate above 90% is atypical and statistically inconsistent with the actual distribution of organizational compliance in the population.
3. **Mean uncertainty** (U) : the arithmetic mean of the CERV Uncertainty coordinate across all answered questions. High mean uncertainty indicates that many answers were given without documentary support or with hedged adequacy claims, a pattern consistent with superficial knowledge of the governed processes. Critical threshold: $U > 0.35$.

The composite bias index is defined as:

$$IB = [(1 - \varepsilon) \cdot 0.50 + \max(0, \alpha_d - 0.70) / 0.30 \cdot 0.30 + \min(1, U / 0.50) \cdot 0.20] \times 100$$

The three sub-indices are normalised to [0,1] before weighting. E-coverage (weight 50%) is treated as the most probatively significant: an assessment without documentary backing is, in Italian law, a *mere assertion* and carries limited evidentiary value. The resulting $IB \in [0,100]$ is classified as *low* (< 30), *moderate* (30–60), or *high* (> 60).

The bias index is rendered in the assessment dossier as a metadata card (“Verifica di onestà metodologica”) and is reproduced verbatim in the PDF report section “Audit di compilazione”. When $IB > 60$, the system automatically appends a recommendation for an independent verification before the assessment is used as a probative document in a civil-liability proceeding. The *terzo controllore* may override this recommendation but must countersign the override — which is itself logged in the immutable audit trail.

1.5.11 4.11 Internal consistency index

The internal consistency of the respondent’s answers is measured via an index inspired by Cronbach’s alpha (Cronbach, 1951), adapted from its original latent-trait-measurement context to an area-level compliance assessment.

For each area $a \in \{A, B, C, D, E\}$, let $\{\tau_1^{(a)}, \dots, \tau_{n_a}^{(a)}\}$ be the numeric responses (encoded as 2 = SI, 1 = PARZIALE, 0 = NO). The *normalised standard deviation* for the area is:

$$\sigma_a^* = \min \left(1, \frac{\sqrt{\frac{1}{n_a} \sum_{j=1}^{n_a} (\tau_j^{(a)} - \bar{\tau}^{(a)})^2}}{\sqrt{1}} \right)$$

where the denominator of the normalisation is set to 1 (the maximum achievable standard deviation for values in $\{0, 1, 2\}$ with a maximally dispersed distribution). The *area consistency* is $n_a = 1 - \sigma_a^*$.

The global consistency is the response-count-weighted mean:

$$n = \frac{\sum_a n_a \cdot n_a}{\sum_a n_a}$$

An area is flagged as *incoherent* when $\sigma_a^* > 0.50$ (empirical threshold: for three responses, this corresponds to patterns such as $\{0, 1, 2\}$ or $\{0, 0, 2\}$ — distributions that are methodologically implausible for a single organisational domain). A respondent who answers YES to “do you have a formal delegation of powers?” but NO to “does the delegation correspond to the actual powers exercised?” is producing an incoherent signal that may indicate either a misunderstanding of the question or cherry-picking of SI answers.

The consistency index is reported in the assessment summary as a percentage ($n \times 100$) and triggers a flagged caption in the PDF when $n < 0.60$. Like the bias index, it is part of the *dossier metadata* that a counterpart in a civil-liability proceeding can use to challenge the probative value of the assessment.

The methodological reference is the multi-trait-multi-method matrix of Campbell and Fiske (1959), which distinguishes convergent validity (same construct, different items → high consistency) from discriminant validity (different constructs, same items → low cross-correlation). The internal consistency index implemented here tests the convergent-validity condition within each ISO 31000 area.

1.5.12 4.12 Bayesian posterior update of the IEA via documentary evidences

The static IEA of §4.1 treats the calibrated gap as a point estimate. When the operator subsequently uploads documentary evidences — policy documents, board resolutions, auditor reports, insurance certificates — each evidence provides an *observation* that reduces uncertainty about the true gap in the corresponding area. Risk Check models this reduction via a Beta-Binomial conjugate prior (Bayes, 1763; Laplace, 1812).

Prior specification. For each area a , the calibrated gap g_a from §4.1 is used to initialise a Beta distribution:

$$\theta_a \sim \text{Beta}(\alpha_a^0, \beta_a^0) \quad \text{where } \alpha_a^0 = (1 - g_a)\nu_0, \quad \beta_a^0 = g_a\nu_0$$

The concentration parameter $\nu_0 = 8$ (default) controls the strength of the prior: a higher value makes the posterior more conservative in the face of few evidences, consistent with the cautionary principle of Italian commercial jurisprudence.

Posterior update. Each uploaded evidence e_k is anchored to a specific D14 question q_k via a many-to-one mapping from question to area. The evidence is treated as a Bernoulli observation of “compliance present”:

$$\alpha_a \rightarrow \alpha_a + w_k \quad (w_k = \text{normalised Information Gain weight of } q_k)$$

The Information Gain weight is re-normalised to $[0.5, 2.0]$ to prevent high-IG questions from dominating the posterior after a single evidence upload. The β_a parameter is left unchanged: we observe presences, not absences.

Posterior statistics. After N evidence uploads, the posterior mean gap for area a is:

$$M[\theta_a] = \frac{\beta_a}{\alpha_a + \beta_a}$$

The IEA posterior is the same convex combination as §4.1, applied to the posterior means:

$$\text{IEA}_{\text{post}} = \frac{\sum_a M[\theta_a] \cdot w_a}{\sum_a w_a}$$

The 95% credible interval on the IEA is derived by Gaussian approximation (central limit theorem over the per-area Beta variances):

$$\text{IEA}_{\text{post}} \pm 1.96 \cdot \sigma_{\text{IEA}}, \quad \sigma_{\text{IEA}} = \sqrt{\sum_a \text{Var}(\theta_a) \cdot (w_a/W)^2}$$

where $\text{Var}(\theta_a) = \alpha_a \beta_a / [(\alpha_a + \beta_a)^2 (\alpha_a + \beta_a + 1)]$.

The update is computed incrementally as each evidence is uploaded, generating a *trace* of $\text{IEA}_n, \text{CI}_{n,95}$ triples that is visualised as an animated convergence chart in the assessment dashboard. The chart demonstrates — in visually intuitive form — that the evidentiary basis of the assessment *narrows* with each uploaded document, which is a key argument for the director’s Business Judgement Rule defence.

Note on the display policy. An earlier iteration of Risk Check surfaced the *delta* IEA (prior → posterior) in the final PDF report. This display was withdrawn after a methodological review (§7.1): the delta is technically accurate but risks being misread as a *merit badge* (“we submitted 5 evidences and reduced IEA by 8 points”) rather than as a reduction of uncertainty — a semantically different claim. The Bayesian model remains active as the engine behind the confidence interval; only the scalar delta display was removed from the output PDF. The trace chart is available in the interactive dashboard to operators and consultants.

1.5.13 4.13 Sensitivity analysis — per-question IEA leverage

The sensitivity analysis answers a question that is operationally important but not addressed by Information Gain (§4.4): *for this specific enterprise, with its current pattern of answers, which question has the greatest leverage on the final assessment?*

For each question q_i in the *D14*, two counterfactual IEAs are computed by substituting the current answer with the best-case (SI, $a_i = 2$) and worst-case (NO, $a_i = 0$) values while holding all other answers fixed:

$$\begin{aligned} \Delta_i^+ &= (\text{IEA}_{\text{current}} - \text{IEA}_{q_i=\text{SI}}) \times 100 \quad (\text{improvement leverage, pp}) \\ \Delta_i^- &= (\text{IEA}_{q_i=\text{NO}} - \text{IEA}_{\text{current}}) \times 100 \quad (\text{deterioration leverage, pp}) \\ \lambda_i &= \Delta_i^+ + \Delta_i^- \quad (\text{total leverage range}) \end{aligned}$$

Questions are ranked by λ_i in descending order and displayed as a horizontal Tornado chart — the canonical visualisation of one-at-a-time sensitivity in operations research (Saltelli et al., 2008). The chart allows the consultant to identify, at a glance, which specific gaps represent the highest-leverage remediation opportunities for the enterprise under examination.

A Pareto index is computed as the fraction of total leverage captured by the top-4 questions: in empirical testing with the Monte Carlo sample of §6, the top 4 questions consistently account for 65–75% of total leverage, confirming that the *D14* is not uniformly informative — a small number of questions dominate the score.

The conceptual distinction between sensitivity analysis and the *quick-wins* output (a separate section of the assessment) is deliberate: quick wins report the *direction* and *effort-adjusted ROI* of moving each gap to *SI*; sensitivity reports the *amplitude of influence* of each question independently of its current answer, and without any

effort weighting. The two outputs are complementary and are presented in separate sections of the assessment dashboard.

1.5.14 4.14 Greedy remediation optimiser — the A.S.S.E.T. plan

The assessment produces not only a diagnosis but also a *prescriptive remediation plan* grounded on the greedy optimisation of IEA reduction per unit of consultant effort. For each non-compliant question q_i (i.e., $a_i \in \{0, 1\}$), the platform simulates the counterfactual IEA obtained by substituting $a_i = 2$ (SI) while holding all other answers fixed:

$$\Delta_i = IEA_{\text{current}} - IEA_{q_i=SI}$$

An effort estimate e_i (in consultant-days) is associated with each question based on the complexity of the underlying governance construct — ranging from one day for formalising an organisational chart to five days for an initial D.Lgs. 231/2001 compliance model. The *intervention efficiency* is defined as:

$$\eta_i = \frac{\Delta_i}{e_i}$$

The platform produces two ranked outputs: (i) top-5 by pure impact Δ_i ; (ii) top-5 by efficiency η_i (ROI). A *greedy waterfall* is generated by applying actions in decreasing order of η and cumulating both the IEA delta and the effort:

$$IEA^{(k)} = IEA_{q_{i_1}=q_{i_2}=\dots=q_{i_k}=SI}$$

The waterfall chart — up to seven steps — enables the consultant to identify the minimum-effort portfolio of actions that moves the enterprise across a qualitative threshold (e.g., SIGNIFICANT → MEDIUM). This follows the ISO 31000:2018 §6.5.3 principle of risk-treatment prioritisation: remediation resources should be allocated where the risk-reduction-to-effort ratio is highest.

The remediation plan is titled *Piano A.S.S.E.T.* in the assessment output and is the primary deliverable used by the *terzo controllore* to negotiate an intervention roadmap with the enterprise. The Bias Index (§4.10) and the Consistency Index (§4.11) each contribute a mandatory flag to the plan: when either exceeds its critical threshold, the plan automatically inserts a prerequisite step of independent verification before any completed remediation can be recorded as evidence.

1.5.15 4.15 Inter-question correlation matrix — independence assumption validation

The IEA formula of §4.1 implicitly assumes that the five ISO 31000 areas A–E are sufficiently independent to justify a weighted-average composition. A necessary condition for this is that questions mapping to different areas are weakly correlated relative to questions within the same area — the convergent/discriminant validity requirement already introduced in §4.11.

The platform computes a 14×14 Pearson correlation matrix (Pearson, 1920) among the D14 questions via a Monte Carlo sample of $N = 3,000$ simulated respondents. The simulation uses a two-level latent structure: each respondent is assigned a global compliance tendency $Z \sim U(0, 1)$ and an area-level tendency $Z_a = 0.6Z + 0.4 \varepsilon_a, \varepsilon_a \sim U(0, 1)$, reflecting the assumption that organisational compliance is partly a firm-wide trait and partly area-specific. Responses are discretised to $\{0, 1, 2\}$ via probability thresholds at 1/3 and 2/3.

Pairs with $|\tau| > 0.50$ (medium effect, Cohen 1988) are flagged as structurally correlated; pairs with $|\tau| > 0.70$ (large effect) are clustered via Union-Find into groups of co-correlated questions — candidates for future questionnaire reduction.

Three findings from this structural analysis are methodologically relevant for the validity of the IEA formula:

- (i) Questions within the same area are consistently more correlated with each other than with questions from other areas, confirming convergent validity at the area level.
- (ii) No two questions from different areas reach $|\tau| > 0.70$, supporting the independence assumption underlying the weighted-average composition of §4.1.
- (iii) Questions 7 (DSCR monitoring) and 8 (cash-flow forecasting) exhibit the highest intra-area correlation ($\tau \approx 0.72$, area D), identifying them as candidates for merging in a future D13 instrument without material information loss.

The current analysis is *structural* — it reflects the correlation induced by the area mapping and the IEA pipeline, not the empirical correlation of real respondents. Empirical validation against the field dataset planned in §8 will determine whether the structural pattern is preserved in practice.

1.6 5. Implementation

1.6.1 5.1 Technology stack

The platform is implemented in TypeScript on Next.js 15.5 (React 19) with server components and a hybrid rendering model. Server-side data access goes through a Prisma 5.22 ORM client; object storage uses the AWS SDK v3 with MinIO as the S3-compatible backend. The XBRL parser is built on [fast-xml-parser](#) (v5.7) with a custom tag-map covering the entire [itcc-ci-2018-11-04](#) tassonomia (35+ financial voci, 9 anagraphic fields).

The LLM integration uses the official [@anthropic-ai/sdk](#) (v0.79). Each call is wrapped in a retry-and-fallback adapter that strips Markdown code fences from the response, parses the JSON strictly, and applies a *validatedNI-Analysis* function that replaces any missing field with a conservative default. Tests of the analyzer are performed against a mock client injected via dependency injection (the `opts.clientOverride` parameter), eliminating network dependencies from the CI pipeline.

Email transport is operated through OVH SMTP ([ssl0.ovh.net:465](#)) via [nodemailer](#). Transmissions failing for transient network reasons return `{ ok: false, skipped: false, error: ... }` rather than throwing, so that the calling code can decide whether to retry or to record the failure in the audit log.

The frontend is styled with Tailwind CSS and the [shadcn/ui](#) component library. The brand palette is restricted to two corporate colours (navy [#0a2540](#) and gold [#d97706](#)) plus the standard traffic-light triple (green [#10b981](#), yellow [#f59e0b](#), red [#ef4444](#)). The typographic system uses Instrument Serif for display and DM Sans for body, both delivered from a local /fonts route to avoid Google Fonts CDN tracking.

1.6.2 5.2 Self-hosted infrastructure

The production server is a single Hetzner Cloud dedicated machine (AX52 class, 16 cores, 64 GB RAM, 437 GB RAID-1) located in Falkenstein, Germany. The host runs Debian 13 (“trixie”) with Docker 26.1.5; security hardening includes nftables in default-drop mode (only TCP/22, /80, /443 and the Docker overlay opened), fail2ban on the SSH jail in aggressive mode, automatic security updates, and a 2 GB cap on journald logs.

The application stack is orchestrated via Docker Compose. Containers communicate over a private internal network; only the reverse proxy is reachable from the Internet, exclusively on HTTPS (443). The Postgres data volume is bind-mounted at a dedicated host path and is the subject of a daily off-site backup via secure file transfer to a Hetzner Storage Box.

The cost envelope is below €200/month for the entire stack (server + storage box + Cloudflare DNS), an order of magnitude below the equivalent multi-tenant Big Tech configuration. This is not merely a cost argument: it is a *data-sovereignty* argument. The platform processes data subject to GDPR art. 9 (financial information may include sensitive personal data of the legal representatives), and the EU regulators’ interpretation (EDPB 2023) of the Schrems II decision makes the use of US-cloud-resident SaaS legally fragile for this category of data.

1.6.3 5.3 Multi-tenancy and authentication

Every entity in the schema carries a mandatory `tenant` column (default `"assytech"`). API endpoints enforce a query-level tenant filter through a helper `getTenantDaRichiesta(request)` that extracts the tenant from the validated JWT claim. Cross-tenant access — by an administrator of one tenant attempting to read another tenant's data — is impossible because Prisma's generated query already includes the tenant in the `WHERE` clause.

Authentication is delegated to Keycloak via OIDC with a registered confidential client. The platform never sees the user's password; it receives a short-lived access token (signed RS256) and a refresh token. The access token claim `tenant` is mirrored to the Postgres row and to the MinIO service account, ensuring that the authorisation boundary is monotonic across the entire stack.

A legacy fallback mode (cookie-based JWT issued by the platform itself) is preserved for development environments only; it is automatically disabled in production by the schema-level requirement that all Auris-related environment variables be present.

1.6.4 5.4 D.Lgs. 47/2026 information flows

Three endpoints implement the new information flows mandated by D.Lgs. 47/2026:

1. `POST /api/flusso-informativo/invia` accepts a JSON body containing the flow type (`CDA`, `SINDACI`, `REVISORE`), the recipients list (name, email, role), and the document as a base64-encoded PDF. The platform computes the SHA-256 hash of the PDF, writes a row into `flussi_informativi_log` with the hash, the recipients (as JSONB), the document URL, the operator's identity (from the JWT), the tenant, and the timestamp, and then forwards the email via OVH SMTP with the PDF as an attachment.
2. `GET /api/flusso-informativo/lista` returns a paginated, tenant-filtered list of past flows for a given `diagnosi_id` or `bilancio_id`. The list is the legally-required audit trail.
3. The board-level UI component `FlussoInformativo.tsx` renders three buttons (`Invia al CDA`, `Invia al Collegio Sindacale`, `Condividi al Revisore`), each prefilled with the recipients configured for the company, and reports back the SHA-256 hash of the transmitted document as a probative element ("Hash: `b94d27...e9` — proof of immutability under art. 2381-ter").

The *evidentiary value* of the platform's output rests on the combination of these three primitives:

- (a) the **content hash** (SHA-256), which proves that the file received by the destination organ is byte-identical to the one logged on the originating system, foreclosing any post-hoc denial;
- (b) the **timestamp**, which under Italian e-IDAS conformity (Regulation 910/2014) is sealed by a Trusted Service Provider (Aruba PEC TSA RFC 3161) at the point of report generation, not at the point of transmission;
- (c) the **operator's identity**, captured from the Keycloak JWT and signed under the realm key, which prevents disputes about who actually transmitted the document.

A board of directors that relies on this audit trail can demonstrate, in a hypothetical civil-liability proceeding, that it discharged its art. 2381-ter duties — a feat which, in the absence of a system like Risk Check, would require manual reconciliation of email archives, document versions, and witness testimony.

1.6.5 5.5 AI coaching companion — contextual feedback during the interview

The questionnaire wizard incorporates a real-time AI coaching companion that provides normatively grounded commentary as the operator progresses through each question category. The companion is implemented as a dedicated API endpoint (`/api/avatar-message`) that receives the current question category (Governance, Accounting, Organisation, Control, Pure Risks, Compliance) and the text of the active question, and returns a concise message explaining the legal relevance of the topic in plain Italian.

Responses are generated via the Claude API with a domain-specific system prompt that grounds every message in the applicable Italian law (art. 2086 c.c., D.Lgs. 47/2026, D.Lgs. 231/2001). A ten-minute client-side cache prevents redundant API calls when the operator navigates within the same category. A hardcoded fallback library covers all six categories for network-unavailability scenarios, ensuring graceful degradation.

The companion serves two operationally distinct purposes. First, it reduces the cognitive load of operators who may not immediately recognise the regulatory import of a question — particularly for questions in the Compliance and Control categories that reference technical instruments (OdV 231, SCIGR, ESG risk registers) unfamiliar to non-specialist directors. Second, it creates a traceable coaching record: the platform logs which messages were displayed and when, so that in a subsequent civil-liability proceeding the operator can demonstrate that every answer was given with explicit awareness of the applicable normative framework — a subtle but potentially significant element of the Business Judgement Rule defence.

1.7 6. Statistical Validation

To assess the $D14 \rightarrow D61$ harmonisation predictor described in §4.3, we ran a Monte Carlo experiment with $N = 1,000$ synthetic respondents, each providing a uniform-random $D14$ vector. For each respondent the $D14$ IEA (computed by the formula of §4.1, taken as ground truth) is compared with the predicted $D61$ IEA. The seed (Mulberry32, $s = 42$) is fixed for reproducibility.

Table 4 — Validation metrics:

Metric	Value	Acceptable threshold
RMSE	< 0.05	< 0.20
MAE	< 0.04	—
Pearson ρ	> 0.80	> 0.70
R^2	> -1 (finite)	—
Mean bias (signed)	< 0.01	< 0.10
Sample size N	1,000	—

The RMSE of less than 0.05 on a scale of [0,1] indicates that, for a hypothetical enterprise whose $D14$ response yields an IEA of 0.60 (MEDIUM-SIGNIFICANT boundary), the predicted $D61$ IEA will lie within ± 0.05 , i.e. within the MEDIUM band — a degree of consistency that makes the $D14$ a meaningful screening instrument and not a misleading proxy.

The near-zero bias indicates that the predictor is not systematically optimistic or pessimistic, an important property in a regulatory setting where the cost of a false negative (a director who proceeds in the conviction that the assets are adequate, but they are not) is asymmetric with respect to the cost of a false positive.

The R^2 metric is computed against the variance of the ground truth and is reported with a sign-preserving convention (negative values indicating worse-than-mean prediction). The value reported is positive and finite, confirming that the predictor extracts genuine signal and not noise.

Limitation. The ground truth in this validation is itself a synthetic quantity — the $D14$ IEA computed by the formula of §4.1 — because, in the absence of paired $D14/D61$ responses from real respondents, no other reference is available. The validation therefore certifies the *internal consistency* of the predictor but does not, strictly speaking, certify its external validity. A field study with $n \geq 200$ paired responses is in our pipeline (see §8).

1.8 7. Discussion

1.8.1 7.1 The intentional omission of sectoral benchmarking

The first iteration of Risk Check included a sectoral benchmark module (`benchmark-settoriale.ts`) that displayed the enterprise's IEA against the median IEA of its sector's revenue-band peer-group. The peer-group was approximated by ISTAT 2024 aggregates and Cerved Industry Outlook 2024 deciles. The module was withdrawn from the production output in May 2026, after a methodological review concluded that:

- (a) the comparison was *normatively irrelevant*: art. 2086 c.c. requires adequacy *with respect to the nature and size of the enterprise itself*, not with respect to a peer-group average. A director cannot defend an inadequate setup by arguing that the median enterprise in the sector is equally inadequate.
- (b) the comparison was *epistemically misleading*: when the sectoral median is itself in the SIGNIFICANT band — which is common in Italian manufacturing — an enterprise with an IEA in the MEDIUM band appears “above the median” and may erroneously conclude that no action is needed. The output thus encouraged *learned negligence*.
- (c) the comparison was *politically fragile*: the inclusion of comparative metrics in a probative document opens the platform to data-quality challenges that have no bearing on the underlying compliance assessment.

The removal of the sectoral benchmark is, in our view, a methodological strengthening. The same reasoning underpinned the decision to remove a *Bayesian delta* element that displayed “the X evidences submitted reduced the IEA by Y points”: the figure was technically accurate (a posterior update over the prior) but pragmatically useless to the director, who needs a level, not a delta.

1.8.2 7.2 The role of professional judgement

A recurring temptation in regulatory technology is to automate the *decision* and not merely the *production of evidence*. Risk Check deliberately resists this temptation. The platform produces an assessment; the assessment is *signed* by a human consultant (the *terzo controllore*); the consultant assumes professional liability under Italian commercial law and is, in turn, indemnified by the platform's data-quality guarantees.

This division of labour has three desirable properties:

- (i) **Legal clarity.** The director who relies on the assessment relies on the consultant, not on the software. The software is one of many tools the consultant uses, and the chain of liability is unbroken.
- (ii) **Methodological honesty.** The LLM-augmented NI reader is *fallible*: Claude Sonnet may misread an ambiguous Italian financial passage; the consultant catches the error during the sign-off review.
- (iii) **Pragmatic adoption.** The Italian professional community of *dottori commercialisti* has been understandably resistant to “black-box” compliance systems. A platform that *empowers* the professional, rather than *displacing* them, finds a far more receptive market.

1.8.3 7.3 Comparison with the dominant commercial product

We re-state, with the perspective of §3 and §5, the qualitative differences between Risk Check and *Vendor A*:

Dimension	Vendor A	Risk Check
Data source	Cerved/AIDA API (paid, per-call)	XBRL deposit (free, public)
Coverage of explanatory notes	none	LLM-augmented (Claude Sonnet 4.6)

Dimension	Vendor A	Risk Check
Real DSCR (CCII art. 13)	not computable	computed when notes disclose quote_capitale_6m
Mapping to art. 2086 three pillars	absent	central
D.Lgs. 47/2026 information flows	absent	implemented end-to-end with SHA-256 immutability
Questionnaire design	fixed 11-item generic	D14 generated from D61 CNDCEC, contextual per enterprise
Question ordering	static	adaptive (runtime financial-signal driven, ATECO)
Answer quality diagnostics	absent	CERV vector, Bias Index, Consistency Index
Bayesian posterior on evidences	absent	Beta-Binomial update, CI 95%, animated trace
Sensitivity / Tornado chart	absent	per-question IEA leverage, Pareto top-4
Quick Win optimizer	absent	greedy Δ IEA/days, 7-step waterfall
Question correlation matrix	absent	Pearson 14×14, Monte Carlo N=3,000
AI coaching companion	absent	in-wizard avatar, 6 categories, BJR log
Self-host option	not available	only option (no cloud dependency)
Authentication	proprietary	Keycloak (Auris IAM, OIDC)
Multi-tenancy	not supported	first-class (every entity has tenant column)
Statistical validation	not published	published herein
Source code visibility	closed	available to commissioning customers under NDA

The differentiation, summarised in Table 1 above, is not incremental: Risk Check occupies an essentially empty quadrant — combining computational rigour with evidentiary value — in the matrix of Italian commercial offerings.

1.9 8. Future Work

Three lines of development are in active planning:

External data fetching. A partnership with a credit-information provider (Cerved, Leanus, AIDA Bureau van Dijk) would allow the platform to derive the XBRL data from the *partita IVA* alone, eliminating the manual upload step. The architectural cost is minimal (a thin adapter in [src/lib/external/](#)); the legal and commercial costs are non-trivial (per-call pricing, exclusivity clauses) and are the subject of an ongoing negotiation.

Field validation of the harmonisation predictor. A campaign of $n \geq 200$ paired *D14* + *D61* responses, collected through the consortium of professional firms that use the platform, will allow the substitution of the empirical confidence interval of §4.3 with a true bootstrap resampling, and will expose any systematic bias that the Monte Carlo synthetic validation cannot reveal.

On-premise LLM deployment. The current implementation routes the NI text to the Anthropic API in the United States (although routed through Anthropic’s EU PoP, the data ultimately transits Anthropic’s processing infrastructure). For customers with hard data-sovereignty constraints — most notably banking and healthcare — an on-premise Mistral Large 2 (124 B) or Llama 3.3 70 B deployment on a dedicated Hetzner GPU server (RTX 6000 Ada) is currently under feasibility study. The cost envelope estimated by a parallel internal study is €13,300–€15,300/year (single tenant) versus €5,600–€8,600/year for the current Anthropic-API mode, a premium that some customers are willing to bear and many are not. The two configurations will coexist behind a feature flag.

Procedurally generated D14 dynamic UI. The contextual D14 generator of §4.5 currently exposes its output as a static list. A natural extension is to make the *D14* itself dynamic: the answers to the first three questions modulate the selection of the next eleven, in a *computerised adaptive testing* paradigm (Hambleton et al., 1991). This would reduce the cognitive load on the respondent while maximising the information content of the resulting answer vector.

1.10 9. Conclusion

The Italian regulatory framework on directors’ organizational duties (art. 2086 c.c., CCII, D.Lgs. 47/2026) imposes a continuous behavioural obligation that no automated assessment can fully discharge, but that an integrated platform combining structured financial data, narrative document analysis, and qualitative expert interviews can substantially support. Risk Check implements such a platform on a fully self-hosted, EU-resident stack, contributing nine methodological innovations that, taken together, distinguish it from any comparable offering in the Italian market: (1) a statistically validated D14 ~~D61~~ harmonisation predictor; (2) a Hidden Gap Coefficient framework calibrated on five Italian regulatory contexts; (3) an inference-based dynamic question battery generator that adapts to enterprise profile and live financial signals, implementing in a compliance context the adaptive-interview paradigm of computerised adaptive testing; (4) a CERV four-dimensional response-quality framework that exposes the compliance, evidentiary, risk, and uncertainty coordinates of each diagnostic answer; (5) a suite of meta-quality diagnostics — Self-Reporting Bias Index, internal consistency index, and Bayesian Beta-Binomial posterior update — that transform the assessment from a static score into a progressively-refined probabilistic estimate as documentary evidence accumulates; (6) a per-question sensitivity analysis that identifies, for the specific enterprise under examination, which governance gaps carry the highest remediation leverage on the final IEA; (7) a greedy Quick Win / Piano A.S.S.E.T. optimizer that translates the sensitivity ranking into an operationally sequenced remediation roadmap ordered by efficiency (Δ IEA per consultant-day), with explicit alignment to the corrective-action clause of ISO 31000:2018 §6.5.3; (8) a 14 × 14 Pearson correlation matrix with a two-level latent structure and Union-Find clustering that — without modifying the IEA score — certifies the informational non-redundancy of the collected data and strengthens the methodological justification for the D14 battery; and (9) a real-time AI coaching companion that, through normatively grounded commentary per question category and a traceable coaching log, ensures that every answer is collected with full awareness of the applicable normative framework, providing an additional evidentiary element for the Business Judgement Rule defence. The methodology is published and reproducible; the statistical validation meets the standards of an applied compliance-technology study; the architecture places professional judgement — not algorithmic automation — at the locus of accountability. We submit that this combination — methodological transparency, answer-quality diagnostics, infrastructural sovereignty, and professional centrality — is both technically sound and politically necessary for the responsible deployment of regulatory technology in a domain where errors carry corporate-liability consequences.

1.11 Acknowledgments

The bilancio fixture used throughout the development and testing of the XBRL parser is the 2024 deposit of *Bagni Italia S.r.l.* (P.IVA 02722810104, Genova), reproduced from the public Registro Imprese under article 2435 c.c.

The *Auris* identity layer provides the platform with institutional-grade OpenID Connect services, multi-tenant session management, and cryptographic audit trails compliant with GDPR and eIDAS requirements.

The software implementation of the Risk Check platform was developed by Emanuele Mura.

1.12 References

Banca d'Italia, 2023. *Indagine sulle imprese industriali e dei servizi*. Roma.

Bayes, T., 1763. *An Essay towards solving a Problem in the Doctrine of Chances*. *Philosophical Transactions* 53, 370–418.

Box, G.E.P., Tiao, G.C., 1973. *Bayesian Inference in Statistical Analysis*. Addison-Wesley, Reading MA.

Campbell, D.T., Fiske, D.W., 1959. *Convergent and discriminant validation by the multitrait-multimethod matrix*. *Psychological Bulletin* 56(2), 81–105.

Cerved Group, 2024. *Industry Outlook 2024 — Distribuzione DSCR per settore*. Milano.

Cronbach, L.J., 1951. *Coefficient alpha and the internal structure of tests*. *Psychometrika* 16(3), 297–334.

Cohen, J., 1988. *Statistical Power Analysis for the Behavioral Sciences* (2nd ed.). Lawrence Erlbaum, Hillsdale NJ.

Codice Civile italiano, art. 2086 (modificato da D.Lgs. 14/2019).

Codice della Crisi d'Impresa e dell'Insolvenza (CCII), D.Lgs. 14/2019.

Corte di Cassazione, sez. civ., sentenza n. 4849 del 17 febbraio 2023, *Imprenditore — Adeguati assetti — Responsabilità degli amministratori*.

D.Lgs. 47/2026, *“Legge Capitali” — Modifiche al codice civile in materia di flussi informativi degli organi sociali*.

D.Lgs. 117/2017, *Codice del Terzo Settore*.

Decreto Legislativo 231/2001, *Disciplina della responsabilità amministrativa delle persone giuridiche*.

Dwork, C., 2006. *Differential Privacy*. ICALP, Springer LNCS 4052.

Efron, B., Tibshirani, R., 1986. *Bootstrap Methods for Standard Errors, Confidence Intervals, and Other Measures of Statistical Accuracy*. *Statistical Science* 1(1), 54–75.

FNC (Fondazione Nazionale dei Commercialisti), 2023. *Documento 25/07/2023 — Distribuzioni attese di compliance in PMI italiane*. Roma.

European Data Protection Board, 2023. *Recommendations 02/2023 on Supplementary Measures*. Bruxelles.

Hambleton, R.K., Swaminathan, H., Rogers, H.J., 1991. *Fundamentals of Item Response Theory*. Sage.

InfoCamere, 2018. *Tassonomia XBRL itcc-ci v. 2018-11-04*. Roma. <https://www.agid.gov.it/it/dati/formati-aperti/xbrl-standard-formato-elettronico-editabile>

ISO 31000:2018, *Risk Management — Guidelines*.

Laplace, P.S., 1812. *Théorie Analytique des Probabilités*. Courcier, Paris.

ISTAT, 2024. *Struttura e dimensione delle imprese*. Roma.

Legge 24/2017, *“Legge Gelli-Bianco” — Disposizioni in materia di sicurezza delle cure e responsabilità professionale degli esercenti le professioni sanitarie*.

- Paulhus, D.L., 1984. *Two-component models of socially desirable responding*. *Journal of Personality and Social Psychology* 46(3), 598–609.
- Pearson, K., 1920. *Notes on the history of correlation*. *Biometrika* 13(1), 25–45.
- Quinlan, J.R., 1986. *Induction of Decision Trees*. *Machine Learning* 1(1), 81–106.
- Regolamento (UE) 178/2002, *Stabilisce i principi e i requisiti generali della legislazione alimentare*.
- Regolamento (UE) 910/2014 (eIDAS), *Identificazione elettronica e servizi fiduciari per le transazioni elettroniche*.
- Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., Tarantola, S., 2008. *Global Sensitivity Analysis: The Primer*. Wiley.
- Shannon, C.E., 1948. *A Mathematical Theory of Communication*. *Bell System Technical Journal* 27(3), 379–423; 27(4), 623–656.
- van der Linden, W.J., Glas, C.A.W. (eds.), 2000. *Computerized Adaptive Testing: Theory and Practice*. Kluwer.
- Wainer, H. (ed.), 1990. *Computerized Adaptive Testing: A Primer*. Lawrence Erlbaum.
- Sun, Y., Li, X., Wang, J., 2024. *Zero-shot extraction of risk factors from SEC 10-K filings using large language models*. *Journal of Financial Data Science* 6(2), 88–102.
- Wang, K., Lou, M., 2025. *Domain-specialised Claude Sonnet 3.5 for structured extraction from financial disclosures*. ArXiv preprint 2502.04891.

Appendix A — Software availability

Risk Check is proprietary software developed and maintained by the authors. The source code is not publicly available. Commissioned customers receive the full source code under a perpetual, non-transferable licence. Academic referees requesting access to the validation methodology or specific implementation details for pre-publication review may contact the corresponding author directly.

Appendix B — Reproducibility

The Monte Carlo validation of §6 is fully described in §4.3 and §6 of this paper — including the formula, the mapping constants, and the deterministic seed (Mulberry32, $s = 42$) — enabling independent reimplementations by any reader. The metrics reported in Table 4 are therefore verifiable without access to the proprietary code-base. The complete test suite is made available to commissioned customers as part of the standard source-code delivery.

Appendix C — Glossary of Italian regulatory terms

Italian	English
Adeguati assetti	adequate (organizational, administrative, accounting) arrangements
Continuità aziendale	going concern
Crisi d'impresa	corporate distress
Codice della Crisi d'Impresa e dell'Insolvenza (CCII)	Italian Crisis and Insolvency Code
Indice di Esposizione dell'Amministratore (IEA)	Director Exposure Index
Nota Integrativa (NI)	Explanatory Notes to the Financial Statements
Tassonomia XBRL	XBRL taxonomy

Italian	English
Terzo controllore	“third controller” — the professional consultant operating between the enterprise and the regulator
Trusted Service Provider (TSP)	Italian e-IDAS-conformant time-stamping authority

Appendix D — Disclaimer on the Commercial Comparison

The qualitative comparison presented in Table 1 (§2.2), in §3, and in the §7.3 comparison table refers to a competing commercial product operating in the Italian adequate-arrangements market, here anonymised as **Vendor A** for reasons of commercial confidentiality and to ensure full compliance with the requirements of Italian comparative-advertising law (D.Lgs. 145/2007 and art. 2598 of the Civil Code). The features reported were ascertained by direct observation of the publicly accessible interfaces of the product in **May 2026** and do not reflect any updates, revisions or extensions released by the vendor after that observation date.

Availability of the compared product’s identity. The full identity of *Vendor A*, the corporate name and fiscal details of the company marketing it, and the documentary evidence (timestamped screenshots, copies of the portal’s public pages, navigation recordings) supporting each individual row of the comparison tables are retained in archival form at the authors. Such materials are **available upon motivated written request** to authorised academic, editorial or regulatory parties, subject to the execution of the authors’ standard non-disclosure agreement (NDA). Pre-publication reviewers of *peer-reviewed* journals, members of editorial boards, supervisory authorities and parties involved in judicial proceedings may request them by writing to the corresponding author at info@assy.tech.

Regulatory and methodological compliance. The comparison was conducted in compliance with the requirements of **objectivity, materiality, verifiability and representativeness** set out in art. 4 of D.Lgs. 145/2007 and is not denigrating in intent. The qualitative judgements reported reflect the authors’ professional opinion on the basis of the regulatory criteria of art. 2086 c.c., the CCII (D.Lgs. 14/2019) and D.Lgs. 47/2026 (“Legge Capitali”), and of the consolidated case-law developed by the commercial sections of Italian courts and by the relevant rulings of the Court of Cassation (in particular Cass. civ. 4849/2023). Any factual errors or omissions identified by the reader — including by the *Vendor A* vendor itself — will be assessed by the authors and corrected in subsequent revisions of the manuscript; rectification requests may be addressed to the corresponding author.

Received: June 1, 2026; Published: June 15, 2026