ARTIFICIAL INTELLIGENCE AND MERIT-BASED LEADERSHIP SELECTION IN ROMANIAN STATE-OWNED ENTERPRISES: A DECISION-SUPPORT FRAMEWORK

Mihnea-Tudor Orjan*
Bucharest University of Economic Studies, Bucharest, Romania

Abstract
This article presents a proof of concept for supporting leadership selection in Romanian state-owned enterprises with artificial intelligence while preserving lawful human responsibility. It combines literature review, legal analysis and institutional comparison with an artificial *ex ante* evaluation based on selection procedures from a district-heating operator and a mining company. Sixty-nine source criteria are organised into thirteen merit dimensions, with legal eligibility assessed separately. Sixteen deliberately constructed candidate profiles test sector fit, transferable experience and legal or evidential barriers; they do not represent the executive labour market. The framework produces documented eligibility outcomes, transparent scores and a pool of candidates meeting predefined robustness conditions, not an appointment decision. Sensitivity tests vary weights, thresholds, criterion mappings and individual ratings. The pools remain stable across the tested thresholds and mappings, but small rating changes can alter leading positions and pool membership. The contribution is an auditable process linking evidence, merit and human review within the applicable appointment rules. The study does not establish predictive validity, reduced patronage or better real-world appointments; these require subsequent evaluation using lawful candidate data.

Keywords
state-owned enterprises; corporate governance; artificial intelligence; merit-based selection; executive recruitment; algorithmic accountability; Romania

JEL Classification
G34; L32; M12; M51; C44

Introduction

State-owned enterprises (SOEs) combine commercial, public-service and strategic responsibilities, which makes leadership quality important to both financial stewardship and public value. Romania has formalised appointments through Government Emergency

* Corresponding author**, Mihnea-Tudor Orjan** – orjanmihnea22@stud.ase.ro

Ordinance No. 109/2011 and Government Decision No. 639/2023, using expectations letters, profiles, competency matrices, longlists and shortlists. Yet the Organisation for Economic Co-operation and Development (OECD, 2023) recorded that, at the end of 2021, 194 of 271 board positions and 48 of 79 executive positions in the 50 largest Romanian SOEs were provisional. Although these historical data do not describe the situation in 2026, they show that formal criteria alone do not guarantee transparent, evidence-based selection.

This article contributes the *Merit-based, Explainable, Robust, Integrity-centred and Traceable Institutional Process for public-enterprise leadership* (MERIT-IP). This process separates legal eligibility, merit, evidence and process controls. Its operational output combines eligibility findings, documented scores and a robustness-qualified candidate pool for legally constrained human review.

In this article, artificial intelligence (AI) refers to machine learning, natural language processing, large language models and rule-based systems used for lawful sourcing, document processing, consistency checks and explanation. The ranking core remains deterministic and inspectable. These tools may improve information processing, but they may also conceal patronage, reproduce proxy discrimination and encourage automation bias.

Candidate data are governed by the General Data Protection Regulation (GDPR). Recruitment and candidate-evaluation systems may also fall within Annex III of the Artificial Intelligence Act (AI Act).

MERIT-IP therefore applies high-risk safeguards as a design baseline across the appointment tracks examined.

The theoretical mechanism is a multiple-principal agency chain. Within this chain, public owners, boards, commissions, the Agency for Monitoring and Evaluation of Performance of Public Enterprises (AMEPIP), consultants and technical providers hold different information, powers and incentives. Version-locked criteria, evidence traceability and documented review of system outputs are intended to reduce selected information asymmetries while preserving the appointment powers of the legally competent body. However, the standard and the platform may themselves be subject to second-order capture.

Methodologically, the study is framed as a design-science-informed artefact evaluation. Official procedures of the Bucharest Municipal District Heating Joint-Stock Company (CMTEB) and CUPRU MIN Joint-Stock Company (CUPRU MIN) provide 69 source criteria, which are applied to 16 manually designed stress profiles. The present stage is an artificial *ex ante* evaluation that tests internal logic, traceability and sensitivity without affecting any person.

Naturalistic evaluation is a separate subsequent stage requiring a shadow pilot with lawfully obtained dossiers, independent evaluators and observed organisational outcomes. Artificial evaluation is not treated as a substitute for that stage (Venable, Pries-Heje and Baskerville, 2016).

Appointment authority does not necessarily imply freedom to disregard a ranking. For competitive board appointments in state-owned companies, Article 29(1) of Government Emergency Ordinance No. 109/2011 requires appointment in shortlist ranking order (Government Emergency Ordinance No. 109/2011). Executive appointments follow Article 35(4), (8) and (11): the board of directors or, in two-tier companies, the

supervisory board appoints from the shortlist, but Article 35(8) does not itself expressly prescribe ranking order. Any scope for choice must therefore also be assessed against the approved selection procedure. In the CUPRU MIN case, that procedure expressly requires appointment in ranking order (CUPRU MIN, 2025a, paras. 67 and 69(i)). MERIT-IP therefore treats its robust pool as an analytical output, not as the statutory shortlist or an independent authorisation to depart from it. Where ranking is binding, the system supports evidence processing, scoring checks, sensitivity analysis and effective review rather than alternative appointment discretion.

The study addresses three research questions:

- How can digital support separate legal eligibility, evidence and merit without displacing the applicable appointment rules?
- How sensitive are the resulting rankings and robust pools to weights, thresholds, criterion mappings and human ratings?
- Which human responsibilities and review safeguards are necessary for the framework to provide practical institutional value?

The first question is addressed by the legal-process design in Sections 2.5–2.7, the second by the sensitivity results in Sections 3.1–3.3 and the third by the governance and implementation analysis in Sections 3.4–3.8.

1. Review of the scientific literature

1.1. Politicisation, public mission and the meaning of merit

State ownership gives elected authorities a legitimate role in defining ownership objectives, but not in treating board or executive positions as patronage. The OECD therefore recommends structured, merit-based and transparent nomination processes protected from undue political influence (OECD, 2024). Romanian research also identifies difficulties in managerial-performance assessment and weaknesses in non-financial reporting and governance quality (Munteanu, 2018; Dragomir, Dumitru and Feleagă, 2022). Comparative evidence associates political connections with weaker outcomes in public utilities, while an Italian reform that restricted political appointments improved board quality and several performance measures (Menozzi, Gutiérrez Urtiaga and Vannoni, 2012; Baltrunaite et al., 2025). Formal credentials alone are not enough, because generic experience may transfer poorly across sectors (Sidki, Boerger and Boll, 2024).

Merit is also harder to define in an SOE than in a conventional private firm. Financial results must be considered alongside regulated tariffs, service obligations, safety, accessibility, investment and governance. For example, postponing maintenance may improve short-term profit while undermining the public mandate (Ramamurti, 1987). MERIT-IP therefore treats incompatibilities, conflicts, disqualifications and minimum legal conditions as gates rather than compensable merit criteria. Incomplete evidence instead produces a PENDING status and a request for clarification.

1.2. Structured selection and algorithmic assistance

Selection research favours structured over unbounded human judgement. Pre-specified combination rules, job-related questions, behavioural anchors and independent ratings improve consistency and reliability (Kuncel et al., 2013; Levashina et al., 2014). MERIT-

IP therefore freezes definitions, evidence rules, thresholds and weight ranges before candidate identities are assessed. It treats interviews and statements of intent as evidence rather than repeated merit blocks.

Research on algorithmic recruitment is more cautious. Efficiency and standardisation can coexist with historical bias, proxy discrimination, opacity, reduced autonomy and difficult-to-validate vendor claims (Tambe, Cappelli and Yakubovich, 2019; Köchling and Wehner, 2020; Raghavan et al., 2020). Removing protected fields does not eliminate discrimination through proxies (Barocas and Selbst, 2016). Automation may also weaken personal judgement (Leicht-Deobald et al., 2019). Machine learning can illuminate private-company director selection (Erel et al., 2021), but the historical shareholder-oriented labels used in that setting cannot be transferred directly to the distinct institutional context of Romanian SOEs. Accordingly, AI is confined to auditable support tasks.

Institutional practice points to the same division of labour. The OECD recognises potential uses in sourcing, matching and screening but stresses governance and evidentiary limits, while Support for Improvement in Governance and Management (SIGMA) supports competency frameworks and structured human assessment (OECD, 2025; Lázaro Vicente et al., 2026). These precedents support assisted processing, not autonomous appointment.

1.3. Multiple-principal agency and institutional standardisation

Agency theory begins with delegation under divergent interests and costly monitoring (Jensen and Meckling, 1976). In an SOE, delegation operates through several layers. Public ownership bodies, authorities, boards, executives, AMEPIP, commissions and providers may answer to several government-linked principals with different demands and information (Apriliyanti, Dieleman and Randøy, 2024). Leadership selection is therefore a multiple-principal agency problem rather than a bilateral employer-applicant choice.

MERIT-IP is theorised as a monitoring and contestability device. *Ex ante* criteria, common evidence rules and independent scoring may reduce selected information asymmetries. Legal gates prevent disqualifying facts from being traded against unrelated strengths, sensitivity analysis exposes dependence on modelling choices, and documented review makes corrections and any legally permitted departures traceable. Three propositions follow: traceability should make hidden preference injection more detectable, a robust pool should expose fragile rankings and, where lawful choice remains, increase the cost of unexplained deviation, and central standards require counterweights because methodology, data and suppliers can themselves be captured.

Standardisation may improve comparability, but it also creates distinct institutional and technological risks. Coercive, mimetic and normative pressures can produce organisational homogenisation (DiMaggio and Powell, 1983). Formal procedures may also be adopted ceremonially and decoupled from substantive practice (Meyer and Rowan, 1977). Visible and consequential evaluation criteria may encourage organisations to adapt strategically to the measure rather than to the underlying objective (Espeland and Sauder, 2007). At the candidate level, applicants may likewise tailor their answers or conform their expressed views to perceived selection expectations (Levashina and Campion, 2007). In algorithmic selection, reliance across organisations on the same or

similar systems may create algorithmic monoculture and homogeneous screening outcomes. Proprietary procurement may also create vendor lock-in (Kleinberg and Raghavan, 2021; OECD, 2025; Bommasani et al., 2026). MERIT-IP therefore models legally constrained and contestable algorithmic governance. Standardisation redistributes rather than eliminates discretion and may itself become an object of capture.

1.4. Research gap and contribution

The research gap lies in the limited integration of SOE appointment research, structured selection, algorithmic recruitment and legal-institutional governance. Within the reviewed corpus, no model operationalises their full combination in the Romanian framework: separate legal gates, public-service and sector-specific competencies, a traceable crosswalk, bounded sensitivity analysis, rights safeguards and distinct board and executive appointment routes. This is a qualified corpus-based claim. The contribution is to translate multiple-principal agency and institutional-standardisation mechanisms into auditable design rules, not to assert that no related system exists elsewhere.

2. Research methodology

2.1. Research strategy, evidence and cases

The study combines an integrative literature review, doctrinal legal analysis, institutional comparison and a design-science-informed artefact approach. Romanian and European law is treated as a design constraint. Consistent with the Introduction, the artificial *ex ante* evaluation tests logical consistency, traceability, reproducibility and sensitivity. Deployed utility and predictive validity are reserved for later naturalistic evaluation (Venable, Pries-Heje and Baskerville, 2016). The large number of computational scenarios increases the resolution of sensitivity analysis over a fixed and small synthetic input set. It does not increase the empirical sample size or add external information about real candidates.

CMTEB and CUPRU MIN provide documented, contrasting cases within the Romanian public-enterprise governance framework. Both completed executive-selection procedures provide public profiles, criteria, weights and final scores, allowing the proposed design to be traced back to source material. District heating emphasises continuity of an essential public service, whereas mining highlights sector-specific operational, safety and investment requirements. Their comparison tests whether a common process can preserve these different priorities. The cases are chosen to allow for analytical comparison and traceability, not statistical representativeness (CMTEB, 2024a, 2024b; CUPRU MIN, 2025a, 2025b).

The resulting division of human and technical responsibilities is set out stage by stage in Section 3.4, where it is also compared with a human-only procedure.

2.2. Coding, harmonisation and synthetic assessment

The 69 source items were reduced through directed qualitative content analysis (Hsieh and Shannon, 2005). Each remains verbatim in a traceability crosswalk and is coded by legal nature, evidentiary function, permissible AI support and required human responsibility. Statutory conditions were converted into PASS, PENDING and FAIL gates. Evidence instruments were separated from merit, overlapping constructs were

merged, and sector content was retained in H11-H13. Coding was performed by the author and was not independently replicated. Sixteen manually designed profiles were created after the architecture was fixed, but they were not sampled from real *curricula vitae* or calibrated to the Romanian executive market.

The sixteen profiles provide a compact set of contrasting stress cases rather than a sample intended for population inference. Sector specialists (C01 and C02) are contrasted with cross-sector and internationally experienced executives (C04 and C07) and with an investment-programme director (C12), while the remaining profiles vary in operational, governance, restructuring, digital and public-sector experience. C14 tests the distinction between academic or policy expertise and comparable executive responsibility. C15 and C16 separately test whether high merit can compensate for a legal impediment or incomplete evidence; by design, it cannot. Thus, sixteen profiles provide coverage of these constructed contrasts, not an empirically justified sample size. Their fixed ratings allow the effects of modelling choices to be examined without introducing new candidate data.

2.3. Crosswalk, cut-off and local rating-sensitivity tests

Because no independent recoder was available, no agreement statistic is reported. The sensitivity tests do not substitute for inter-coder reliability and were added to expose how strongly the numerical outputs depend on the author's coding choices. Author-specified plausible remappings are tested in 6,208 configurations through two analyses: a conservative capped-transfer analysis and a deliberately severe source-salience envelope. The difference between the case counts is purely combinatorial. CMTEB generated 6,144 configurations from nine binary and one three-state mapping choices, whereas CUPRU MIN generated 64 from four binary choices after both mapping methods and both transfer intensities were applied. The cases were enumerated and interpreted separately, so the counts are neither statistical sample sizes nor cross-case weights. Nine combinations of the top-four scenario-share cut-off (40%, 50%, 60%) and the lower-tail score floor (70, 75, 80), applied to each case (18 configurations in total), test the demonstrative thresholds. Human-score vulnerability is examined separately by enumerating every admissible one-point edit to the fixed 1-5 ratings: 317 for CMTEB and 320 for CUPRU MIN. These deterministic one-cell edits are local rating-sensitivity tests, not estimates of assessor-error probabilities or an upper bound on coordinated manipulation.

2.4. Weight sensitivity and reproducibility

The central score uses the case-specific weights (table no. 1). To test weight sensitivity, the analysis draws 20,000 independent vectors from a bounded Dirichlet distribution centred on each case vector. The concentration parameter controls how tightly the draws cluster around the central vector, with lower values permitting wider dispersion and higher values keeping them closer to the baseline. The value $\kappa=500$ is an intermediate design benchmark, not an empirically estimated parameter. The values $\kappa=250$ and 1,000 test looser and tighter concentration. Independent batch rejection sampling continues until 20,000 feasible vectors satisfy the published bounds and sum to one. The method is a constrained parametric Monte Carlo sensitivity analysis, not Bayesian inference (Butler, Jia and Dyer, 1997). The reported diagnostics cover convergence, acceptance rates, conditional-mean displacement and Monte Carlo error.

2.5. The MERIT-IP model: institutional and process architecture

Institutional routing precedes calculation. For board appointments, the public owner, the selection and nomination commission, and AMEPIP act under Government Emergency Ordinance No. 109/2011, while the competent shareholder body makes the appointment. For executive appointments, the board of directors or supervisory board retains appointment authority under company law, the special framework and the constitutive act (Law No. 31/1990; Government Emergency Ordinance No. 109/2011; Government Decision No. 639/2023; OECD, 2024). One engine may support both tracks, but it cannot transfer decision rights. The process consists of vacancy monitoring, a version-locked profile, lawful sourcing, dossier confirmation, legal review, structured assessment, aggregation, explanation and reconsideration.

2.6. Formal decision rules

Let candidate i be assessed against legal requirements m and merit dimensions k. The overall gate is:

$$G_i = \text{FAIL if any } g_{im} = \text{FAIL; PENDING if none fails and any } g_{im} = \text{PENDING; PASS otherwise} \quad (1)$$

Only PASS candidates enter the ranking. Equation (2) uses fixed central weights and ratings from 1 to 5, and multiplication by 20 converts the weighted average into an index ranging from 20 to 100. The five critical dimensions H3, H4, H9, H11 and H12 require a rating of at least 3/5. Outside those threshold-protected dimensions, the additive model permits full compensation between criteria. This design was chosen for transparency and auditability, not because compensation is normatively inevitable. Future work should compare less-compensatory or outranking specifications where stakeholders reject particular trade-offs.

$$S_i^0 = 20\,\Sigma_k\, w_k^0 s_{ik}; \quad T_i = 1\{\min_{k \in C}\, s_{ik} \geq 3\}, \quad C = \{H3,H4,H9,H11,H12\} \quad (2)$$

Ratings remain fixed in the baseline demonstration. In scenario b, only the weight vector varies under equation (3), where κ is the concentration parameter and ℓ_k and u_k are the admissible bounds. Because conditioning changes marginal means, these means are reported as diagnostics rather than reinterpreted as central weights.

$$w^{(b)} \sim \text{Dirichlet}(\kappa w^0) \mid \ell_k \leq w_k^{(b)} \leq u_k, \quad \Sigma_k w_k^{(b)} = 1 \quad (3)$$

Equation (4) defines rank-acceptability shares and robust-pool membership. The lower-tail condition $Q_{0.10}(S_i) \geq 75$ means that at least 90% of sampled admissible-weight scenarios score at or above 75. The cut-offs are transparent demonstration parameters, not statutory or empirically validated thresholds.

$$b_i^{(r)} = B^{-1}\Sigma_\beta 1\{\text{rank}_i^{(b)} = r\}; \quad R_i = 1\{G_i = \text{PASS}, T_i = 1, \Sigma_{r=1}^{4} b_i^{(r)} \geq 0.50, Q_{0.10}(S_i) \geq 75\} \quad (4)$$

The cut-offs implement a demonstrative balance between comparative standing and minimum merit. The 3/5 thresholds protect governance, safety, ethics, sector obligations and operational continuity against compensation by unrelated strengths. Top-four placement in at least half of the admissible-weight scenarios retains several repeatedly competitive alternatives, while the lower-tail floor of 75 guards against low scores under less favourable weights. Top-four is a scenario-ranking condition, not a legally prescribed shortlist size or a requirement that the final pool contain exactly four candidates. These values are transparent design choices rather than validated selection standards; the alternative cut-offs test dependence on them rather than establish their optimality.

The harmonised criteria and alternative weights are used here for artificial *ex ante* evaluation. They do not retrospectively replace the approved matrices of either case. In an operational procedure, the official score must follow the lawfully adopted and published rules. Thus, sensitivity scenarios remain separate diagnostic outputs. Where the official ranking is binding, robust-pool membership and first-place shares remain such outputs and do not create an alternative legally operative ranking.

2.7. Legality by design and meaningful human control

AI may search approved sources, structure evidence, flag inconsistencies, calculate scenarios and prepare explanations. It may not infer political opinions or protected traits, decide incompatibilities, reject a candidate irreversibly or make the appointment. Accessibility and reasonable accommodation are required throughout (Government Ordinance No. 137/2000; Law No. 202/2002; Law No. 448/2006). GDPR principles and effective review apply. *SCHUFA Holding* treats a score on which a decision strongly relies as potentially subject to Article 22, while *Dun & Bradstreet Austria* requires intelligible information on the procedure used (OQ v Land Hessen, 2023; CK v Magistrat der Stadt Wien, 2025). Where the intended purpose falls within Annex III, AI Act Articles 8-15 require risk management, documentation, logging, human oversight, accuracy and cybersecurity from 2 December 2027 (Regulation (EU) 2016/679, 2016; Regulation (EU) 2024/1689, 2024; Regulation (EU) 2026/1744, 2026).

Human judgement enters before calculation through legal qualification, evidence assessment and criterion ratings, and after calculation through the reasoned appointment decision. Aggregation and sensitivity analysis between those stages are mechanical and reproducible. Meaningful human control requires evaluators to score independently before seeing the aggregate and to inspect the evidence and its limitations. They must also have the authority to alter or stop the output, give comparative reasons, and make correction and reconsideration available (Article 29 Data Protection Working Party, 2018). Robust-pool membership has no independent legal effect. Where the applicable rules make the ranking binding, the official ranking governs, and additional reasons cannot cure a departure from it; enhanced reasons for choosing outside the pool become relevant only where the applicable rules leave lawful choice. FAIL cannot be overridden by managerial preference, and PENDING must be resolved within the applicable clarification rules and deadlines before ranking.

Table no. 1. Harmonised criteria and case-specific central weights

Identifier	Criterion	CMTEB / CUPRU MIN	Admissible range	Minimum
H1	Strategy and mandate	8 / 8%	6-10%	-
H2	Financial stewardship	8 / 8%	6-10%	-
H3	Governance/reporting/performance	8 / 8%	6-10%	3/5
H4	Risk, safety and resilience	8 / 9%	5-9%	3/5
H5	People and social dialogue	7 / 6%	5-9%	-

Identifier	Criterion	CMTEB / CUPRU MIN	Admissible range	Minimum
H6	Digital/organisational change	5 / 5%	3-7%	-
H7	Stakeholders and negotiation	5 / 4%	3-7%	-
H8	Analytical judgement	5 / 5%	3-7%	-
H9	Ethics and independence	10 / 10%	8-12%	3/5
H10	Comparable leadership	6 / 7%	5-9%	-
H11	Sector/regulation/public service	10 / 11%	8-12%	3/5
H12	Operations/quality/continuity	13 / 11%	10-14%	3/5
H13	Investment delivery	7 / 8%	6-10%	-

Source: author's design based on the official CMTEB and CUPRU MIN selection matrices and the normative constraints described in the text.

3. Results and discussion

3.1. Harmonisation and case results

Harmonisation changed the logic of the source matrices before ranking began. CUPRU MIN placed conflicts and incompatibilities alongside merit even though a verified impediment was exclusionary, so MERIT-IP converts them into gates. In CMTEB, twelve statement-of-intent items represented 41.34% of the raw weighting and overlapped with strategy, objectives, risk and investment, so MERIT-IP treats them as a transversal work sample. Reputation, independence and integrity together represented about 0.10% of the normalised weight. MERIT-IP assigns ethical leadership a substantive weight and threshold (CMTEB, 2024a; CUPRU MIN, 2025a). Being integrity-centred does not mean that H9 must carry the largest numerical weight. Ethics is protected by both a substantive 10% weight and a non-compensatory threshold, while operational continuity in both cases, and sector competence in CUPRU MIN, carry slightly larger case-specific weights because continuity, safety and sector fit are core public obligations.

At κ=500, the CMTEB robust pool is {C01, C04, C07, C12}, with C04 ranking first in every sampled admissible-weight scenario. The CUPRU MIN pool is {C02, C04, C07, C12}. C02 has the highest central score (84.72) and ranks first in 54.87% of scenarios, compared with 45.13% for C04. C15 would lead both scoreboards but is excluded by FAIL, while C16 remains unranked until the PENDING evidence is clarified (table no. 2).

These are properties of constructed stress cases, not findings about actual managers or other real persons.

Table no. 2. Combined proof-of-concept results

Case	Candidate and synthetic profile	Central score	First-place scenario share (%)	Classification
CMTEB	C04 - Infrastructure transformation	85.86	100.00%	Robust

Case	Candidate and synthetic profile	Central score	First-place scenario share (%)	Classification
CMTEB	C07 - International energy	84.38	0.00%	Robust
CMTEB	C01 - District-heating	84.30	0.00%	Robust
CMTEB	C12 - Investment programmes	81.20	0.00%	Robust
CUPRU MIN	C02 - Mining/metallurgy	84.72	54.87%	Robust
CUPRU MIN	C04 - Infrastructure transformation	84.52	45.13%	Robust
CUPRU MIN	C07 - International energy	83.32	0.00%	Robust
CUPRU MIN	C12 - Investment programmes	82.90	0.00%	Robust
Both	C15 - High merit; incompatibility	86.96 / 88.42	Not ranked	FAIL
Both	C16 - High merit; incomplete file	84.78 / 85.02	Not ranked	PENDING

Source: MERIT-IP fixed synthetic ratings and reproducible bounded-Dirichlet simulations (κ=500; 20,000 scenarios).
Note: Percentages are first-place scenario shares, not population probabilities. FAIL and PENDING cases are not ranked.

3.2. Sensitivity, convergence and design-level verification

Weight sensitivity changes the leading candidate in CUPRU MIN but not in either robust pool. C04 leads in 50.95% of scenarios at κ=250, whereas C02 leads in 54.87% at κ=500 and 61.66% at κ=1,000. Rank stability depends on the structure of criterion advantages, not only on the central margin: C04's 1.48-point CMTEB lead rests mainly on the highly weighted H12 and H11, whereas the 0.20-point CUPRU MIN margin reflects offsetting advantages and changes more readily as admissible weights vary. At κ=500, the approximate 95% Monte Carlo error interval for C02's first-place share is [54.18%, 55.56%]; it measures numerical simulation error, not uncertainty about real candidates. Acceptance rates are 18.10% and 10.93%, and maximum conditional-mean weight shifts are 0.53 and 0.95 percentage points for CMTEB and CUPRU MIN, respectively. Estimates stabilise by 20,000 scenarios. The CUPRU MIN 5th–95th percentile range of the C02-C04 score difference spans zero, confirming sensitivity to admissible weights.

3.3. Mapping, cut-off and local rating sensitivity

The robust pools remain unchanged in all 18 cut-off configurations and all 6,208 author-specified crosswalk configurations. Under the conservative remapping analysis, both leaders remain unchanged. Under the severe envelope, the CUPRU MIN order reverses

in 2 of 16 configurations, while pool membership remains stable. Because no independent recoder was used, this exercise tests numerical dependence on coding choices but not inter-coder reliability or construct validity.

Single-rating edits can change who leads or which candidates belong to the robust pool. Of the 317 admissible one-point CMTEB edits, 10 change pool membership and 12 change the leader. Of the 320 CUPRU MIN edits, 5 change the pool and 20 change the leader. Taken together, Section 3.2 and the sensitivity results (table no. 3) show that robust-pool stability under the tested changes to weights, cut-offs and remappings does not guarantee stability under changes to individual ratings. In practical terms, the 15 pool-membership and 32 leader changes among the 637 enumerated single-rating edits mean that one disputed rating on a decisive criterion can change who ranks first or which candidates form the robust pool; apparently small scoring disagreements therefore warrant evidence-based review. Coordinated edits may have cumulative effects and were not exhaustively enumerated here. A stable aggregation rule cannot correct unsupported or retrospectively altered inputs. Score provenance, assessor separation, evidence traceability and version control are therefore substantive safeguards rather than administrative details.

Table no. 3. Sensitivity and validation results

Case	Cut-off grid	Mapping cases	Mapping pool changes	Mapping leader changes	±1 rating edits	Rating pool / leader changes
CMTEB	9/9 unchanged	6,144	0	0 / 6,144	317	10 / 12
CUPRU MIN	9/9 unchanged	64	0	4 / 64	320	5 / 20

Source: author's calculations. Crosswalk alternatives are author-specified. Each ±1 test changes one candidate-criterion rating in isolation. These are local one-cell rating-sensitivity tests, not estimates of assessor-error probabilities or upper bounds on coordinated manipulation. The four CUPRU MIN mapping leader changes comprise two of sixteen configurations at each severe-envelope transfer intensity. None occurs in the capped-transfer analysis.

3.4. From ranking to a contestable governance architecture

A point score is a weak governance device unless legal status, evidence and sensitivity remain visible. MERIT-IP is designed to constrain and expose selected capture opportunities, not to abolish politicisation. Criteria are frozen before identities are assessed, versions and delays are logged, common work samples are used, scores are linked to evidence, independent ratings precede aggregation, and corrections require documented reasons. The sensitivity tests qualify this claim. Pool membership is stable to the examined cut-offs and mappings, while local one-cell rating edits can change leaders and, less often, pool membership.

A human-only procedure is compared with the proposed support process (table no. 4). Human-only does not mean unstructured or unauditable: the middle column describes

responsibilities that remain human in both arrangements. The final column identifies the additional technical support proposed by MERIT-IP.

Table no. 4. Human responsibilities and additional MERIT-IP support

Stage	Human-only procedure and retained human responsibility	Additional MERIT-IP support
Criterion definition	Approve lawful criteria, evidence rules and published weights.	Store the approved version and flag subsequent changes.
Legal eligibility	Verify legal facts, decide eligibility and authorise permitted clarifications.	Check completeness, flag inconsistencies and apply human-confirmed gates.
Evidence review	Assess relevance and reliability; resolve evidential gaps.	Organise evidence and link assessments to source material.
Rating	Rate each dimension independently and justify the assessment.	Record ratings without assigning or silently rewriting them; preserve versions.
Aggregation	Ensure correct application of the approved formula and verify the resulting ranking.	Calculate scores reproducibly and perform arithmetic consistency checks.
Sensitivity analysis	Specify and interpret sensitivity tests, where undertaken.	Run predefined scenarios and identify changes in leading positions and pool membership.
Traceability	Retain assessments, amendments and decision records.	Link inputs, rules, scores and amendments through an audit trail.
Explanation and reconsideration	Explain findings, examine objections and authorise corrections.	Prepare traceable summaries and recalculate after authorised corrections.
Final appointment	The competent body appoints under the binding legal and procedural rules.	Provide the supporting record; neither appoint nor authorise departure from binding rules.

Source: author's proposed process design.

Note: This is a design comparison, not an observed comparison of appointment outcomes. The proof of concept evaluates the deterministic scoring and sensitivity core; document-processing assistance and workflow controls require operational validation.

The intended advantages are specific but not yet validated. Version-locking makes later changes to criteria or weights detectable; separating legal gates from merit prevents disqualifying facts from being offset by unrelated strengths; automated consistency checks could flag discrepancies between stated ranks and reported scores, such as those in the CUPRU MIN final report (CUPRU MIN, 2025b, p. 8), without implying manipulation; and one-point tests show which ratings materially affect the outcome and therefore warrant closer evidentiary review. The claim is not that these safeguards are

impossible without the system or that their real-world benefits have already been demonstrated.

Standardisation can itself be gamed. Authorities may tailor the profile, teach candidates to the rubric, restrict sourcing, validate evidence selectively or use the platform as ceremonial legitimation. Centralisation may create a single point of capture at AMEPIP or the supplier. Counterweights include open sourcing, public version histories, assessor separation, red-team tests, supplier portability, audit rights and periodic revalidation.

3.5. Institutional deployment and legally bounded human review

Romania need not create another appointing authority. Within its current remit, AMEPIP could issue guidance, host pilot infrastructure, administer the existing pool of candidates for public-enterprise board appointments, monitor deadlines and require audit-ready outputs where permitted by law. Nationwide mandatory use, new register links or new processing purposes would require express legislation. AMEPIP reported 648 cumulative admissions to the administrator pool and separately consulted the market on recruitment services for about 400 board positions (AMEPIP, 2026a, 2026b). Shareholder bodies must retain board appointments, while boards or supervisory boards retain executive appointments.

Human control must operate within the applicable appointment rules. Where the final ranking is binding, a sensitivity warning calls for documented verification and any legally available correction or reconsideration, while it does not authorise appointment of a lower-ranked candidate. Where lawful choice remains, the decision must stay within the legally admissible candidate set and follow the published criteria. Enhanced reasons and independent review for choosing outside the model's robust pool are proposed safeguards only in that setting, not a general legal power to override ranking. Candidates should receive intelligible reasons and access to the applicable review process.

3.6. Rating-stage safeguards against error and capture

The single-rating tests in Section 3.3 show that a one-point change can alter the leading candidate or robust-pool membership. Rating controls should therefore be a priority in any shadow or controlled pilot. At least two trained evaluators should rate each dimension independently before seeing aggregate results. Irrelevant identity information should be withheld where feasible, although detailed career histories prevent guaranteed anonymity. Candidate order may be randomised. Disagreements beyond a predeclared threshold should trigger documented reconciliation, preserving original ratings. Scores and amendments should be timestamped and version-locked in tamper-evident records, with identifying data linked only through controlled access. Across repeated cycles, rater-level patterns may inform audit where sample sizes permit, but statistical flags must not be treated as proof of error or manipulation. The one-point tests identify sensitivity, not a calibrated model of assessor behaviour. These safeguards concern the production and review of rating inputs, and they do not change Equations (1)–(4).

3.7. Transfer, cost-benefit and technical safeguards

Foreign and private practice offers useful tools, including talent mapping, semantic sourcing, standardised dossiers and workflow automation. It does not, however, provide a ready-made public appointment system (OECD, 2025). Romanian SOEs also require

open access, public-mission criteria, legal gates, reasons, review and vendor accountability. A defensible cost-benefit protocol should measure implementation costs and expected avoided appointment loss. In this protocol, expected avoided appointment loss would be assessed as the change in the risk of a weak appointment multiplied by the economic and public-value gap over the mandate, plus service, safety, litigation and reputational externalities. These quantities are not estimated here and require a non-decisional shadow pilot. The shadow pilot is therefore the immediate empirical next step, not full operational deployment.

Interoperability should rely on purpose-limited interfaces to authorised registers, role-based access, separation of identifying and scoring data, encryption, version locking, tamper-evident logs, incident response and adversarial testing. Procurement must allocate AI Act and GDPR roles, prohibit undisclosed model changes, preserve documentation and portability, regulate subcontractors and permit independent audit. Trade secrecy cannot leave the public authority unable to explain its decision.

3.8. Phased implementation

Implementation should proceed from retrospective testing to a non-decisional shadow pilot and then to a controlled decision-support pilot. Limited scaling should occur only after data-protection, fairness, security, usability and independent-audit checks. Autonomous appointment is excluded from every phase.

Conclusions

MERIT-IP contributes an auditable process for separating legal eligibility, merit assessment, evidence and final appointment responsibility. Its artificial *ex ante* evaluation shows that the robust pools remain stable across the tested weight concentrations, cut-offs and criterion remappings, whereas one-point rating edits can change leading positions and pool membership. The main practical lesson is that transparent aggregation must be accompanied by reliable, independently reviewed scoring inputs.

The framework does not demonstrate better appointments, reduced patronage or predictive validity. Its safeguards address selected information asymmetries, while recognising that criteria, evaluators and technical providers can themselves become targets of capture. The process remains AI-assisted and human-decided: human oversight means responsibility for eligibility findings, ratings, corrections and the lawful appointment act, not unrestricted discretion. The next step is a non-decisional shadow pilot assessing lawful data use, rating reliability, fairness and operational usefulness before deployment. The practical contribution ultimately depends on the binding effect of the applicable final ranking: where that ranking determines appointment, MERIT-IP supports transparent evidence processing, auditable scoring, consistency checks, sensitivity analysis and effective review rather than selection from an alternative robust pool.

The review is integrative rather than PRISMA systematic (*Preferred Reporting Items for Systematic Reviews and Meta-Analyses*: a structured reporting framework for systematic literature reviews), and institutional sources may describe initiatives favourably. Only two executive procedures and manually designed profiles are used. The 69-to-13 crosswalk was coded by one author, while disclosure and remapping tests support traceability but not independent construct validity. Evidence-routing ambiguities cannot

be propagated without item-level candidate scores. Confidential dossiers, assessor-level ratings, performance outcomes and lawful protected-group data were legally unavailable. The model tests logic, traceability and sensitivity, not predictive validity, candidate experience, actual discrimination or reduced patronage. The present evaluation is artificial and *ex ante*, and naturalistic validation is a separate phase that cannot be inferred from the stress tests. Fixed ratings omit assessor severity. Local one-cell edits reveal vulnerability, but they are neither a calibrated error model nor a coordinated-manipulation test. A future shadow pilot should use lawful dossiers, multiple trained evaluators, inter-rater analysis, candidate feedback, proxy and group diagnostics, accommodation outcomes and longitudinal financial, operational, governance, integrity, safety, environmental and public-service indicators.

Further research should independently recode the crosswalk and refine the competency ontology through a multi-round Delphi process involving corporate-governance, human-resources, legal and sector experts. Both consensus and unresolved disagreement should be reported (Hasson, Keeney and McKenna, 2000). Such a process could strengthen content validity and reduce dependence on one author's coding choices without eliminating normative judgement. Future work should also elicit weights formally, optimise collective board composition, estimate avoided appointment loss and test gaming, ceremonial compliance, cybersecurity and supplier concentration.

Data and code availability. The manuscript reports all information required to interpret the analysis. The complete crosswalk, synthetic inputs, executable code, workbooks and detailed diagnostic outputs are archived in Zenodo at https://doi.org/10.5281/zenodo.21739408. Legal and institutional sources were initially verified on 31 July 2026. The provisions and procedural documents concerning appointment authority and the binding effect of rankings were rechecked on 13 September 2026.

Conflict of Interest Statement. The author declares no financial, professional, institutional, personal, or other competing interests in relation to any Romanian or foreign state-owned enterprise, including any entity examined, cited, or otherwise referred to in this article.

References

[1] Agency for Monitoring and Evaluation of Performance of Public Enterprises (AMEPIP), 2026a. *First-quarter report, 2026*. [pdf] AMEPIP. Available at: <https://amepip.gov.ro/wp-content/uploads/2026/04/RAPORT-AMEPIP-Q1-2026_aprilie-2026.pdf> [Accessed 31 July 2026].

[2] Agency for Monitoring and Evaluation of Performance of Public Enterprises (AMEPIP), 2026b. *Market-consultation report*. [pdf] AMEPIP. Available at: <https://amepip.gov.ro/wp-content/uploads/2026/03/Raport-privind-consultarea-pietei.pdf> [Accessed 31 July 2026].

[3] Apriliyanti, I.D., Dieleman, M. and Randøy, T., 2024. Multiple-principal demands and CEO compliance in emerging market state-owned enterprises. *Journal of Management Studies*, [e-journal] 61(6), pp.2406-2436. https://doi.org/10.1111/joms.12977.

[4] Article 29 Data Protection Working Party, 2018. *Guidelines on automated individual decision-making and profiling, WP251rev.01*. Endorsed by the European Data Protection Board. [online] Article 29 Data Protection Working Party. Available at: <https://www.edpb.europa.eu/documents/guideline/automated-decision-making-and-profiling_en> [Accessed 31 July 2026].
[5] Baltrunaite, A., Cannella, M., Mocetti, S. and Roma, G., 2025. Board composition and performance of state-owned enterprises: quasi-experimental evidence. *The Journal of Law, Economics, and Organization*, [e-journal] 41(1), pp.239-263. https://doi.org/10.1093/jleo/ewad014.
[6] Barocas, S. and Selbst, A.D., 2016. Big data's disparate impact. *California Law Review*, [e-journal] 104(3), pp.671-732. https://doi.org/10.15779/Z38BG31.
[7] Bommasani, R., Bana, S.H., Creel, K.A., Jurafsky, D. and Liang, P., 2026. Algorithmic monocultures in hiring. In: Association for Computing Machinery, *Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency*. Montreal, QC, Canada, 25-28 June 2026. New York, NY: Association for Computing Machinery. pp.6351-6382. https://doi.org/10.1145/3805689.3812400.
[8] Bucharest Municipal District Heating Joint-Stock Company (CMTEB), 2024a. *Componenta integrală a planului de selecție pentru Director General și Director Economic* [Full component of the selection plan for the General Director and Economic Director]. [online] CMTEB. Available at: <https://www.cmteb.ro/guvernanta-corporativa.php> [Accessed 31 July 2026].
[9] Bucharest Municipal District Heating Joint-Stock Company (CMTEB), 2024b. *Raport final al procedurii de selecție pentru Director General și Director Economic* [Final report of the selection procedure for the General Director and Economic Director]. [online] CMTEB. Available at: <https://www.cmteb.ro/guvernanta-corporativa.php> [Accessed 31 July 2026].
[10] Butler, J., Jia, J. and Dyer, J., 1997. Simulation techniques for the sensitivity analysis of multi-criteria decision models. *European Journal of Operational Research*, [e-journal] 103(3), pp.531-546. https://doi.org/10.1016/S0377-2217(96)00307-4.
[11] *CK v Magistrat der Stadt Wien (Dun & Bradstreet Austria)* (Case C-203/22) [2025] EU: C:2025:117. [online] Available at: <https://curia.europa.eu/juris/document/document.jsf?docid=295841&doclang=en> [Accessed 31 July 2026].
[12] CUPRU MIN Joint-Stock Company, 2025a. *Planul integral de selecție pentru Director General aprobat 28 martie 2025* [Full selection plan for the General Director, approved on 28 March 2025]. [online] CUPRU MIN. Available at: <https://cuprumin.ro/anunturi-de-selectie/> [Accessed 31 July 2026].
[13] CUPRU MIN Joint-Stock Company, 2025b. *Raportul final al selecției Directorului General* [Final report of the selection of the General Director]. [online] CUPRU MIN. Available at: <https://cuprumin.ro/anunturi-de-selectie/> [Accessed 31 July 2026].
[14] DiMaggio, P.J. and Powell, W.W., 1983. The iron cage revisited: institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, [e-journal] 48(2), pp.147-160. https://doi.org/10.2307/2095101.

[15] Dragomir, V.D., Dumitru, M. and Feleagă, L., 2022. The predictors of non-financial reporting quality in Romanian state-owned enterprises. *Accounting in Europe*, [e-journal] 19(1), pp.110-151. https://doi.org/10.1080/17449480.2021.2018474.
[16] Erel, I., Stern, L.H., Tan, C. and Weisbach, M.S., 2021. Selecting directors using machine learning. *The Review of Financial Studies*, [e-journal] 34(7), pp.3226-3264. https://doi.org/10.1093/rfs/hhab050.
[17] Espeland, W.N. and Sauder, M., 2007. Rankings and reactivity: how public measures recreate social worlds. *American Journal of Sociology*, [e-journal] 113(1), pp.1-40. https://doi.org/10.1086/517897.
[18] *Government Decision No. 639/2023 approving the methodological rules for Government Emergency Ordinance No. 109/2011, consolidated version*. [online] Available at: <https://legislatie.just.ro/Public/DetaliiDocumentAfis/273382> [Accessed 31 July 2026].
[19] *Government Emergency Ordinance No. 109/2011 on the corporate governance of public enterprises, consolidated version*. [online] Available at: <https://legislatie.just.ro/Public/DetaliiDocument/133685> [Accessed 31 July 2026].
[20] *Government Ordinance No. 137/2000 on discrimination, consolidated version*. [online] Available at: <https://legislatie.just.ro/Public/DetaliiDocument/24129> [Accessed 31 July 2026].
[21] Hasson, F., Keeney, S. and McKenna, H., 2000. Research guidelines for the Delphi survey technique. *Journal of Advanced Nursing*, [e-journal] 32(4), pp.1008-1015. https://doi.org/10.1046/j.1365-2648.2000.t01-1-01567.x.
[22] Hsieh, H.-F. and Shannon, S.E., 2005. Three approaches to qualitative content analysis. *Qualitative Health Research*, [e-journal] 15(9), pp.1277-1288. https://doi.org/10.1177/1049732305276687.
[23] Jensen, M.C. and Meckling, W.H., 1976. Theory of the firm: managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, [e-journal] 3(4), pp.305-360. https://doi.org/10.1016/0304-405X(76)90026-X.
[24] Kleinberg, J. and Raghavan, M., 2021. Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences*, [e-journal] 118(22), e2018340118. https://doi.org/10.1073/pnas.2018340118.
[25] Köchling, A. and Wehner, M.C., 2020. Discriminated by an algorithm: a systematic review of discrimination and fairness by algorithmic decision-making in HR recruitment and development. *Business Research*, [e-journal] 13(3), pp.795-848. https://doi.org/10.1007/s40685-020-00134-w.
[26] Kuncel, N.R., Klieger, D.M., Connelly, B.S. and Ones, D.S., 2013. Mechanical versus clinical data combination in selection and admissions decisions: a meta-analysis. *Journal of Applied Psychology*, [e-journal] 98(6), pp.1060-1072. https://doi.org/10.1037/a0034156.
[27] *Law No. 202/2002 on equal opportunities, consolidated version*. [online] Available at: <https://legislatie.just.ro/Public/DetaliiDocument/35778> [Accessed 31 July 2026].
[28] *Law No. 31/1990 on companies, consolidated version*. [online] Available at: <https://legislatie.just.ro/Public/DetaliiDocument/798> [Accessed 31 July 2026].

[29] *Law No. 448/2006 on the rights of persons with disabilities, consolidated version*. [online] Available at: <https://legislatie.just.ro/Public/DetaliiDocument/77815> [Accessed 31 July 2026].

[30] Lázaro Vicente, B., Zielinski, W., Salgado Velo, J.F. and Gray, Á., 2026. *Merit-based civil service recruitment and selection in the Western Balkans and EU Eastern Neighbourhood: evidence, OECD practices and reform directions*. SIGMA Papers, 79. [pdf] Paris: OECD Publishing. Available at: <https://doi.org/10.1787/dc514448-en> [Accessed 20 October 2026].

[31] Leicht-Deobald, U., Busch, T., Schank, C., Weibel, A., Schafheitle, S., Wildhaber, I. and Kasper, G., 2019. The challenges of algorithm-based HR decision-making for personal integrity. *Journal of Business Ethics*, [e-journal] 160(2), pp.377-392. https://doi.org/10.1007/s10551-019-04204-w.

[32] Levashina, J. and Campion, M.A., 2007. Measuring faking in the employment interview: development and validation of an interview faking behavior scale. *Journal of Applied Psychology*, [e-journal] 92(6), pp.1638-1656. https://doi.org/10.1037/0021-9010.92.6.1638.

[33] Levashina, J., Hartwell, C.J., Morgeson, F.P. and Campion, M.A., 2014. The structured employment interview: narrative and quantitative review of the research literature. *Personnel Psychology*, [e-journal] 67(1), pp.241-293. https://doi.org/10.1111/peps.12052.

[34] Menozzi, A., Gutiérrez Urtiaga, M. and Vannoni, D., 2012. Board composition, political connections, and performance in state-owned enterprises. *Industrial and Corporate Change*, [e-journal] 21(3), pp.671-698. https://doi.org/10.1093/icc/dtr055.

[35] Meyer, J.W. and Rowan, B., 1977. Institutionalized organizations: formal structure as myth and ceremony. *American Journal of Sociology*, [e-journal] 83(2), pp.340-363. https://doi.org/10.1086/226550.

[36] Munteanu, I., 2018. The challenges of performance assessment in Romanian state-owned enterprises. In: C. Brătianu, A. Zbuchea and A. Vitelar, eds. *Strategica: Challenging the Status Quo in Management and Economics*. Bucharest, Romania, 11-12 October 2018. Bucharest: Tritonic. pp.1247-1259.

[37] *OQ v Land Hessen (SCHUFA Holding (Scoring))* (Case C-634/21) [2023] EU:C:2023:957. [online] Available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62021CJ0634> [Accessed 31 July 2026].

[38] Organisation for Economic Co-operation and Development (OECD), 2023. *OECD review of the corporate governance of state-owned enterprises in Romania*. [e-book] Paris: OECD Publishing. Available at: <https://doi.org/10.1787/fabf20a8-en> [Accessed 20 October 2026].

[39] Organisation for Economic Co-operation and Development (OECD), 2024. *OECD Guidelines on Corporate Governance of State-Owned Enterprises 2024*. [e-book] Paris: OECD Publishing. Available at: <https://doi.org/10.1787/18a24f43-en> [Accessed 20 October 2026].

[40] Organisation for Economic Co-operation and Development (OECD), 2025. *Governing with artificial intelligence: the state of play and way forward in core government functions*. [e-book] Paris: OECD Publishing. Available at: <https://doi.org/10.1787/795de142-en> [Accessed 20 October 2026].

[41] Raghavan, M., Barocas, S., Kleinberg, J. and Levy, K., 2020. Mitigating bias in algorithmic hiring: evaluating claims and practices. In: Association for Computing Machinery, *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. Barcelona, Spain, 27-30 January 2020. New York, NY: Association for Computing Machinery. pp.469-481. https://doi.org/10.1145/3351095.3372828.

[42] Ramamurti, R., 1987. Performance evaluation of state-owned enterprises in theory and practice. *Management Science*, [e-journal] 33(7), pp.876-893. https://doi.org/10.1287/mnsc.33.7.876.

[43] Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation), 2016. *Official Journal*, [online] L119, 4 May, pp.1-88. Available at: <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng> [Accessed 31 July 2026].

[44] Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act), 2024. *Official Journal*, [online] L, 2024/1689, 12 July. Available at: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng> [Accessed 31 July 2026].

[45] Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence, 2026. *Official Journal*, [online] L, 2026/1744, 24 July. Available at: <https://eur-lex.europa.eu/eli/reg/2026/1744/oj/eng> [Accessed 31 July 2026].

[46] Sidki, M., Boerger, L. and Boll, D., 2024. The effect of board members' education and experience on the financial performance of German state-owned enterprises. *Journal of Management & Governance*, [e-journal] 28(2), pp.445-482. https://doi.org/10.1007/s10997-022-09663-4.

[47] Tambe, P., Cappelli, P. and Yakubovich, V., 2019. Artificial intelligence in human resources management: challenges and a path forward. *California Management Review*, [e-journal] 61(4), pp.15-42. https://doi.org/10.1177/0008125619867910.

[48] Venable, J., Pries-Heje, J. and Baskerville, R., 2016. FEDS: a framework for evaluation in design science research. *European Journal of Information Systems*, [e-journal] 25(1), pp.77-89. https://doi.org/10.1057/ejis.2014.36.