Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

The Governance Frontier COBIT-Aligned AI Decision Support and Risk-Adaptive Optimization for Enterprise Technology Modernization

Audrey Rah

arahimi@uh.edu

University of Houston https://orcid.org/0000-0001-7501-7809

David.w Johnson

University of Houston

Systematic Review

Keywords: AI governance, decision support, distributionally robust optimization, enterprise modernization, executive decisionmaking, human oversight, IT governance, multi-expert systems, systematic literature review

Posted Date: November 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11784311/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

The Governance Frontier: COBIT-Aligned AI Decision Support and Risk-Adaptive Optimization for Enterprise Technology Modernization

Audrey Rah[1,*], David W. Johnson[2]
[1]Department of Electrical and Computer Engineering, University of Houston, Houston, TX, USA
[2]Assistant Vice President, University of Houston, UIT Technology Services and Support, Houston, TX, USA
[*]Corresponding author: arahimi@uh.edu David W. Johnson: dwjohnson@uh.edu
ORCID: Audrey Rah – 0000-0001-7501-7809

***Abstract*—Executives deciding enterprise modernization, legacy-system transformation, AI adoption, organizational restructuring, digital-infrastructure change, and cross-functional governance must weigh legal, financial, cybersecurity, procurement, operational, and AI-governance questions that sit in separate functions. The central methods of this paper are GAIM-DF (Governance-Aware Institutional Modernization Decision Framework), a human-gated advisory process for these decisions, and Risk-Adaptive AI Governance Optimization (RA-GO v2), a risk-adaptive optimizer with mandatory governance gates, halt-safe feasibility, distributionally robust evaluation, and preference-information triage for scarce human review; an accountable human decides. As an optional architectural extension, we computationally evaluate a multi-expert coordination layer that convenes specialist capabilities by context and consolidates constraints conservatively. Evidence is staged across four complementary pre-deployment approaches: analytical verification; retrospective historical-data analysis; comparative computational benchmarking against strong baselines; and scenario-based stress testing under distributional shift. Together they provide complementary pre-deployment evidence for high-stakes governance analysis. A systematic review of university and public-sector governance, with observational process-log analyses and RA-GO scenario stress tests, yields six principal findings: role coordination is associated with shorter approval and process times under supporting falsification checks; robust evaluation preserves tail-risk protection under distribution shift; preference-information triage concentrates human review where it most reduces modeled regret; and mandatory gates, an intensity floor, and a halt state are jointly consistent. The multi-expert benchmarks report accuracy, false-positive, latency, and token-cost trade-offs and routing failure modes relative to a strong Single-Agent baseline and to RA-GO alone. Shadow-mode piloting and a risk-adaptive Hybrid remain future work.**

***Index Terms*—AI governance, decision support, distributionally robust optimization, enterprise modernization, executive decision-making, human oversight, IT governance, multi-expert systems, systematic literature review.**

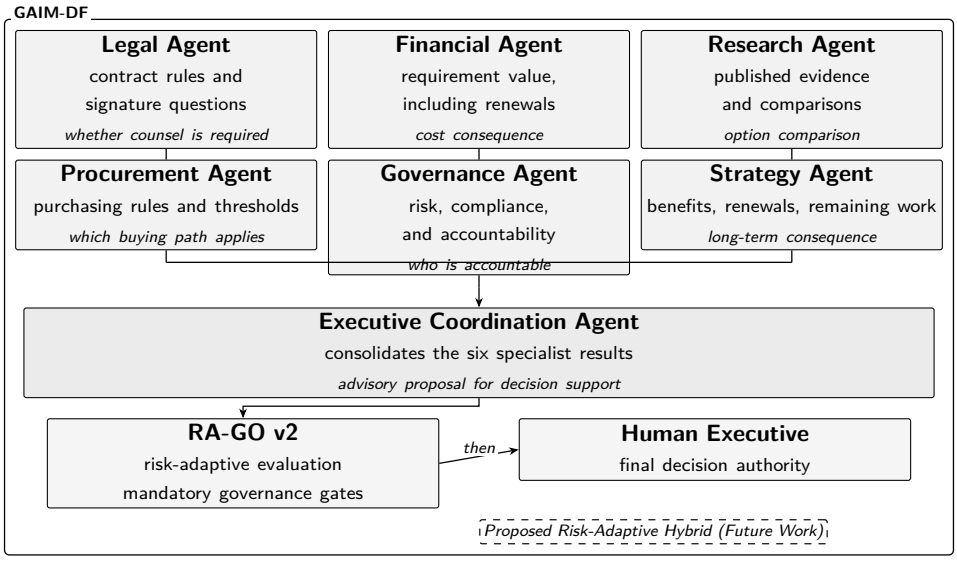


Fig. 1. Illustrative GAIM-DF / RA-GO v2 decision-support view with six specialist roles, executive coordination, mandatory governance gates, and final human decision authority. The optional multi-expert layer evaluated in Section VII instead uses ten triggerable capabilities as logical functions (Table XIII), not a one-to-one mapping to these six roles. The Hybrid extension is proposed for future evaluation.

I. Introduction

EXECUTIVES who decide enterprise modernization—AI adoption, legacy-system transformation, digital-infrastructure change, organizational restructuring, and cross-functional governance redesign—decide with expertise split across legal, financial, cybersecurity, procurement, operational, and AI-governance functions, and are rarely expert in more than one or two. Organizations respond with serialized approval chains. Each step usually protects a legitimate interest, but together the steps can lengthen decisions, fragment ownership, and encourage workarounds, and they were not built for AI-era executive decisions, where an AI component adds privacy, bias, explainability, and oversight questions. Oversight bodies document the cost of aging systems: a 2025 U.S. Government Accountability Office (GAO) review of 69 federal legacy systems identified 11 as most in need of modernization, and only three had modernization plans with all three key elements [1]; GAO also lists IT acquisition and management as a high-risk area [2]. These are federal findings. The literature is conditional: centralized coordination can yield standardization and security benefits yet create lags [3], hierarchy preserves legal certainty [4], and resistance reflects both disposition [5] and how change is designed [6]–[9].

This paper's central methodological contributions are GAIM-DF, a human-gated governance decision process, and Risk-Adaptive AI Governance Optimization (RA-GO v2), a risk-adaptive optimizer with mandatory gates; an accountable human retains decision authority. Empirical grounding is a systematic review of higher education and comparable public bodies, together with observational associations and numerical experiments summarized as six findings (Sections III–V, VIII). Separately, an optional multi-expert coordination layer is specified and computationally compared with a strong Single-Agent baseline and with RA-GO alone to characterize coordination trade-offs and failure modes (Section VII; Fig. 8, Table XIII). Figure 1 is an illustrative six-role overview and is distinct from that ten-capability register.

Prior reviews address IT governance in higher education [10]–[12], barriers to digital transformation [13], [14], and AI governance in education and public administration [15]–[17]. Adjacent enterprise-AI decision and responsibility syntheses examine productivity, workforce change, and human accountability at the point of executive choice [18], [19]. Related multi-agent and AI-governance work is surveyed in Section II-E.

A. Scope of the Evidence and the Six Findings

The framework is stated for organizations generally. The review covers higher education and comparable public bodies; the analyses use a university payment log, a permit log, and public records; and the numerical experiments use constructed scenarios. Six findings, each with its evidence class, are summarized in Table I: role coordination is associated with -52.2% managerial approval latency (F1) and -31.7% adjusted process time (F2); three falsification stages are null as predicted (F3); KL-ball robust evaluation preserves tail-risk protection under distribution-shift testing (F4); Dirichlet expected value of preference information targets scarce executive review (F5); and mandatory gates, the intensity floor, and a halt state are jointly consistent (F6). F1–F3 are adjusted associations, F4–F5 are model behavior under assumed parameters, and F6 is an analytic invariant. The multi-expert evaluations are reported separately from F1–F6.

B. Research Questions

Table II lists the research questions (RQs). RQ8 is answered by the comparison in Section VII, not by the systematic review. Evidence for RQ1–RQ6 is predominantly qualitative, mixed, or conceptual, so a narrative synthesis without meta-analysis is used.

C. Contributions

- C1 An uncapped three-index PRISMA 2020 synthesis of 69 full-text-verified, MMAT/JBI-appraised studies (from 13,287 retrieved and 11,265 deduplicated records, with 53 reports not retrieved tracked in a sensitivity log).
- C2 A critical comparison of governance and decision-making approaches and theoretical lenses across full-text-verified higher-education and public-sector studies.
- C3 An analytical crosswalk from barrier categories to COBIT 2019 objective titles (13 full texts apply COBIT) and verified NIST CSF 2.0, AI RMF 1.0, AI 600-1, SP 800-53, and SP 800-161 controls, with human decision authority retained.
- C4 GAIM-DF, a ten-step human-gated advisory framework for executive modernization decisions, with a Gate H1 constraint catalogue distinguishing generally applicable control classes from jurisdiction-specific examples.
- C5 RA-GO v2, extending GAIM-DF with a three-way active/status-quo/HALT partition, KL-ball distributionally robust optimization (DRO), and Dirichlet EVPI_s review triage (F4–F6; Section VIII).
- C6 A governance-frontier model separating protective from deadweight friction (Propositions 1–5), with observational associations on $6,886$ university payment and $7,065$ permit cases (F1–F3).
- C7 Executive guidelines linking GAIM-DF / RA-GO v2 to COBIT 2019 EDM objectives for technology-governance practice.
- C8 An optional multi-expert coordination extension computationally compared with RA-GO alone and strong Single-Agent baselines on accuracy, false-positive rate, latency, token cost, and routing failure modes, together with a Hybrid design proposed for future work (Section VII).

The remainder of this paper is organized as follows. Section II presents the theoretical background, governance standards, and related research on AI-assisted decision support. Section III describes the systematic review methodology, evidence selection, quality appraisal, and computational evaluation design. Section IV synthesizes the literature on administrative modernization, institutional barriers, governance approaches, and AI oversight. Section V discusses the findings and develops the governance-frontier model. Section VI introduces GAIM-DF. Section VII evaluates an optional multi-expert extension and proposes a Hybrid architecture. Section VIII formulates and evaluates RA-GO v2. Section IX provides executive guidelines,

TABLE I
THE SIX PRINCIPAL FINDINGS, THEIR EVIDENCE CLASS, AND HOW EACH BEARS ON THE CENTRAL QUESTION: CAN EXECUTIVE MODERNIZATION DECISIONS BE SHORTENED WITHOUT WEAKENING THE SAFEGUARDS THEY EXIST TO ENFORCE?

ID	Verified result	Evidence class	Section	Bearing on the central question
F1	-52.2% approval latency under consolidated budget-owner/supervisor sign-off ($\hat{\beta} = -3.03$ d, 95% CI $[-3.32, -2.74]$)	Observational association, confounder-adjusted	VIII-G	Managerial approval delay is associated with how sign-off roles are coordinated, after adjustment for unit, time, activity, and spend controls.
F2	-31.7% overall process time after fixed effects and controls ($\hat{\beta} = -3.83$ d, 95% CI $[-4.93, -2.74]$; -34.9% unadjusted)	Observational association, confounder-adjusted	VIII-G	The association appears in the end-to-end process window, not only in an intermediate approval stage.
F3	Three falsification stages null as predicted ($p_{\text{cl}} = 0.24, 0.21, 0.41$)	Statistical robustness check on observational data	VIII-G	Null results in stages where role consolidation has no mechanism reduce the plausibility of a purely unit- or season-wide artifact.
F4	KL-DRO holds $\text{CVaR}_{0.95}$ and cap safety under structural OOD shifts ($\phi_{\text{matched}} = -0.212$ M\$; 0.0% vs. 6.7% and 26.7% cap violations)	Scenario-based stress testing, in- and out-of-distribution	VIII-E	Under distribution-shift testing, robust evaluation continues to protect tail risk and cap safety when the assumed risk model is misspecified.
F5	Dirichlet EVPI_s triage captures $1.56\times$ the stakeholder-regret reduction per human review of random escalation	Analytic result plus numerical experiments	VIII-F	In the model, EVPI_s triage returns more stakeholder-regret reduction per review than random escalation.
F6	Mandatory gates, intensity floor, and HALT are jointly consistent ($0/160$ violations, $88/88$ unviable cases halted, $0/20{,}000$ proposition violations)	Analytic, verified numerically	VIII-D	Gates, the intensity floor, and an explicit halt state can be made jointly consistent, with halting when no admissible action remains.

TABLE II
RESEARCH QUESTIONS AND WHERE THEY ARE ADDRESSED

ID	Research question	Section	Primary output
RQ1	Which governance and decision-making approaches for administrative and enterprise technology modernization are reported in higher education and comparable public bodies, and how strong is the supporting evidence?	IV-C	Table VI
RQ2	Which barriers (structural, political, cultural, technological, financial, capacity, and AI-related) are reported?	IV-D	Fig. 3, Table VII
RQ3	Which theoretical lenses are used to explain resistance and institutional inertia?	IV-E	Fig. 4, Appendix A
RQ4	How can reported barriers be related to recognized control frameworks (COBIT 2019, NIST CSF 2.0, NIST AI RMF 1.0)?	IV-F	Fig. 5, Table VIII
RQ5	Which strategies are reported to reduce friction while preserving safeguards, and with what evidence?	IV-G	Table IX
RQ6	What roles, risks, and oversight requirements are reported for AI-assisted executive decision support?	IV-H	Text
RQ7	What research gaps remain, and what governance-aware decision framework can be derived from the evidence?	V, VI	Table X, Fig. 7, Algorithm 1
RQ8	Under what conditions, trade-offs, and failure modes does an optional multi-expert coordination layer differ from a strong Single-Agent baseline and from RA-GO alone, including under incomplete intake?	VII	Fig. 8, Tables XIII, XIV, XV, XVI, XVII; Figs. 10, 11; Hybrid proposal Fig. 12

Section X discusses limitations, and Section XI concludes. Appendices A–E provide supporting material.

II. BACKGROUND

A. Institutional Resistance and Structural Inertia

Institutional theory explains why organizations in the same field converge on similar structures through coercive, mimetic, and normative pressures [20]. Path dependence stresses self-

reinforcing sequences that make earlier arrangements costly to reverse [21], and structural inertia argues that established organizations change slowly relative to their environments [22]. Universities have also been characterized as loosely coupled systems and as organized anarchies in which problems, solutions, and participants meet irregularly [23], [24]. These lenses predict friction but do not by themselves say whether a given control is wasteful or protective; that judgment requires evidence about the purpose and outcome of the control.

B. Acceptance, Capability, and Public Value Lenses

Technology acceptance models explain individual adoption [25], [26] but are poorly matched to institution-level approval chains, where the unit of analysis is a decision process rather than a user. Dynamic capabilities [27] and digital transformation research [28], [29] address how organizations sense, seize, and reconfigure, while public value governance [30] frames legitimacy and accountability as outcomes alongside efficiency. Sociotechnical systems thinking [31] adds that technical and social arrangements must be designed jointly. Appendix A compares these lenses and records how often the included studies use them.

C. Control Frameworks

COBIT 2019 organizes governance and management objectives into the domains EDM, APO, BAI, DSS, and MEA [32]. The NIST Cybersecurity Framework 2.0 groups outcomes into six functions (Govern, Identify, Protect, Detect, Respond, Recover) [33], and the NIST AI Risk Management Framework 1.0 into four (Govern, Map, Measure, Manage) [34]. The NIST AI RMF Generative AI Profile adds risks specific to generative systems, including confabulation, data privacy, information security, and value-chain integration [35]. NIST SP 800-53 Rev. 5 includes the control families Assessment, Authorization, and Monitoring (CA), Planning (PL), System and Services Acquisition (SA), and Supply Chain Risk Management (SR) [36], and SP 800-161 Rev. 1 addresses cybersecurity supply-chain risk management practices [37]. GAO's AI accountability framework is organized around governance, data, performance, and monitoring [38]. Complementary treatments of IT risk controls, operational-record auditing, and identity-recovery exposure reinforce the same control vocabulary for trust and residual-risk assessment [39]–[41]. These frameworks supply control and risk vocabularies. The crosswalk in Section IV-F and the EDM alignment in Table XI use those titles analytically.

D. Automation and Human Oversight

Levels-of-automation research distinguishes what functions are automated from how much authority is delegated [42]. Management research describes an automation–augmentation paradox in which the same AI capability can substitute for or complement human judgment depending on design [43]. For approval decisions with legal and security consequences, this motivates advisory designs in which AI supports retrieval and drafting while accountable humans decide [19].

E. Multi-Expert AI Decision Support and the Research Gap

A targeted OpenAlex topic scan (outside the PRISMA corpus) situates the architecture relative to adjacent multi-agent and AI-governance research. Language-model agents and role-based collaboration have been studied for software engineering [44], medical reasoning [45], and scientific hypothesis generation [46], with surveys [47]–[50] and benchmarks for business tasks [51], [52]; selective expert activation originates in mixture-of-experts models [53] and model routing [54]. Related architectures address tool use and multi-agent governance [55], [56], retrieval-augmented enterprise prioritization [57], and semantic tracing for LLM agent workflows [58]. AI governance reviews map themes and requirements [59]–[61], and accountability research examines how automated decisions can be reviewed and contested [62]–[64]. Human–AI studies show that people discount algorithms after seeing them err and use imperfect ones more when they can modify them [65]–[68], which supports keeping recommendations advisory [43]. Within this targeted scan, we did not identify a single prior system that combines rule-governed dynamic expert activation, consolidated non-compensatory gates, risk-adaptive robust optimization, and mandatory human authority for executive modernization decisions; the scan is not an exhaustive multi-database claim of global uniqueness.

III. Methodology

A. Design and Reporting

This systematic literature review follows PRISMA 2020 (Preferred Reporting Items for Systematic reviews and Meta-Analyses) [69] and PRISMA-S [70], structured via SPIDER (sample, phenomenon of interest, design, evaluation, research type) [71]. The review was not registered. Human-participant elicitation is outside this design.

B. Information Sources and Uncapped Search

To obtain uncapped, programmatically reproducible retrieval, we executed 14 Boolean searches on 15 November 2026 across three open scholarly APIs (Table III): (i) OpenAlex (`api.openalex.org/works`, S1–S7, retrieving all 6,348 matches via full pagination); (ii) the U.S. Department of Education's official ERIC API (`api.ies.ed.gov/eric/`, E1–E4, peer-reviewed filter, retrieving all 283 2015–2026 matches); and (iii) the Semantic Scholar Academic Graph Bulk API (`api.semanticscholar.org`, SS1–SS3, retrieving all 6,656 matches). Across the 11,265 deduplicated records, cross-publisher indexing captured 187 IEEE (`10.1109/*`), 89 ACM (`10.1145/*`), 377 Springer, 368 Taylor & Francis, 247 Emerald, 184 Elsevier, and 178 Wiley works. The executable archive inventory is listed in Appendix B.

C. Uniform Eligibility Criteria

Table IV specifies the uniform inclusion and exclusion criteria applied to all 11,265 unique records. Every deduplicated record was screened against IC1–IC4 and EC1–EC5. Peer-reviewed conceptual, legal, and design-science governance analyses meeting IC1–IC4 were eligible and were appraised with the corresponding JBI checklist.

TABLE III
EXPANDED UNCAPPED SEARCH DATABASES, QUERIES, DATES, AND VERIFIED RETRIEVED COUNTS (15 NOVEMBER 2026)

ID	Database / API	Boolean query as executed (2015–2026; title and abstract; uncapped retrieval)	Matches	Retrieved
S1	OpenAlex API	("higher education" OR university) AND ("digital transformation" OR "legacy system" OR "IT modernization") AND (bureaucracy OR "organizational resistance" OR "institutional inertia" OR "decision making")	816	816
S2	OpenAlex API	("higher education" OR university) AND ("IT governance" OR "information technology governance" OR COBIT OR "digital governance")	791	791
S3	OpenAlex API	("higher education" OR university) AND ("resistance to change" OR "institutional politics" OR "shared governance") AND (technology OR digital)	602	602
S4	OpenAlex API	("public sector" OR university) AND (procurement OR "technology acquisition" OR "software acquisition") AND ("information technology" OR "information system")	277	277
S5	OpenAlex API	("artificial intelligence" OR "decision support") AND (governance OR accountability OR "human oversight") AND ("higher education" OR "public administration")	3,591	3,591
S6	OpenAlex API	("process mining" OR "multi-agent" OR "robotic process automation") AND ("higher education" OR university OR "public administration") AND (administrative OR governance)	107	107
S7	OpenAlex API	("institutional theory" OR "path dependence" OR "legacy system") AND ("public sector" OR university) AND (modernization OR "digital transformation")	164	164
E1	ERIC	("higher education" OR university OR college) AND ("digital transformation" OR "legacy system" OR "IT modernization") AND (bureaucracy OR resistance OR governance OR "decision making")	62	62
E2	ERIC	("higher education" OR university) AND ("IT governance" OR "information technology governance" OR COBIT OR "technology governance" OR "data governance")	26	26
E3	ERIC	("higher education" OR university) AND ("procurement" OR "software acquisition" OR "technology acquisition" OR "vendor") AND ("information technology" OR "information system" OR software OR cloud)	11	11
E4	ERIC	("higher education" OR university) AND ("artificial intelligence" OR "decision support" OR algorithmic) AND (governance OR accountability OR "human oversight" OR policy) AND (administration OR institutional)	184	184
SS1	Semantic Scholar Bulk Graph API	("higher education" | "university") + ("IT governance" | "information technology governance" | "COBIT") + ("modernization" | "digital transformation" | "decision" | "procurement" | "resistance")	5,872	5,872
SS2	Semantic Scholar Bulk Graph API	("higher education" | "university") + ("legacy system" | "bureaucracy" | "institutional inertia" | "shared governance") + ("information technology" | "digital transformation" | "software")	151	151
SS3	Semantic Scholar Bulk Graph API	("higher education" | "university" | "public administration") + ("artificial intelligence" | "multi-agent") + ("governance" | "human oversight" | "procurement") + ("administrative" | "decision support")	633	633
Total retrieved across 3 open scholarly APIs before deduplication (100% of matches retrieved; uncapped)				13,287

Cross-publisher metadata in the deduplicated corpus ($n = 11,265$) includes 187 IEEE (10.1109/⋆), 89 ACM (10.1145/⋆), 377 Springer, 368 Taylor & Francis, 247 Emerald, 184 Elsevier, and 178 Wiley records.

D. Selection Process and Full-Text Retrieval

From 13,287 raw records, 594 within-database cross-query duplicates and 1,428 cross-database DOI/normalized-title duplicates (2,022 total) were removed, leaving 11,265 unique records (Fig. 2). One reviewer advanced 129 reports to full-text retrieval and performed the subsequent full-text appraisal. For the remaining 11,136 records, title/abstract exclusion reasons were assigned with a documented rule-based classifier (abstract <45 words →EC4a; title pedagogy/clinical/non-comparable-sector regex →EC1/EC3; missing venue or Zenodo/SSRN DOI →EC5; otherwise EC2), yielding EC2 7,469, EC1/EC3 1,133, EC4a 1,525, and EC5 1,009 (screening_ledger_11265.csv).

Full texts were sought through publisher open-access endpoints, PMC/EuropePMC, ERIC, institutional repositories, and Unpaywall-linked locations (66_studies_access_links.pdf; References/ill_fulltext/retrieval_log.csv). 76 reports were assessed at full text, including [72]–[74]. Of these, 7 were subsequently excluded: three because their bodies were non-English (FT1), two pedagogy/education-use studies (FT2), one fictionalized teaching case (FT3), and one ACODE survey whitepaper that fails IC1 (FT4; [75]). The remaining 69 studies were verified English article bodies (≥1,500 extracted words matching the indexed title). 53 of the 129 sought reports were not retrieved as verified English bodies; they are listed in reports_not_retrieved_ill_list.csv and are not counted among the 69 included studies.

E. Extraction and Formal MMAT 2018 / JBI Appraisal

The 69 studies were extracted from verified article body texts (reference lists stripped) and scored item-by-item using MMAT 2018 (Mixed Methods Appraisal Tool) [76] for qualitative, quantitative descriptive/non-randomized, and mixed-methods empirical studies, the 11-item JBI (Joanna Briggs Institute) Systematic Review Checklist for reviews, and the 6-item JBI Text/Legal/Artifact Checklist for doctrinal, policy, and design-science studies (formal_mmat_jbi_appraisal.csv). Following MMAT 2018 guidance against official overall cutoffs [76], item-level responses are primary; percentage bands (≥80%, 60–79%, <60% of applicable items met) are author-defined descriptive summaries.

F. Synthesis and Sensitivity Check

Synthesis followed thematic and framework synthesis [77]–[79] and SWiM reporting principles [80]. Non-exclusive

TABLE IV
UNIFORM INCLUSION AND EXCLUSION CRITERIA APPLIED ACROSS ALL 11,265 DEDUPLICATED RECORDS

Code	Criterion
IC1	Published 2015–2026; peer-reviewed journal or conference article or systematic review; English full text.
IC2	Setting is higher education administration/IT governance, or a public-sector administration comparator reporting transferable evidence on IT/software procurement, legacy modernization, or AI governance.
IC3	Substantively analyzes technology governance, decision rights, bureaucracy, institutional resistance, software/cloud procurement, legacy/ERP modernization, or AI-assisted administrative decisions.
IC4	Full text retrieved and verified ($\geq$1,500 words of English body text) for primary synthesis; reports sought for retrieval but not obtained as verified English full texts, while meeting IC1–IC3 on abstract, are tracked in PRISMA "not retrieved" ($n = 53$) and retained in an abstract-level sensitivity log.
EC1	Pedagogy, classroom teaching, curriculum, student learning outcomes, K–12, or clinical/hospital care focus.
EC2	Narrow operational IT-service/helpdesk/network audit or individual tool acceptance without institutional governance, procurement, or modernization analysis.
EC3	Private-sector, industrial, or municipal setting without administrative IT/AI governance transferability.
EC4	Insufficient abstract (<45 words) at screening; or non-English body text (FT1), pedagogical focus (FT2), or fictionalized classroom case (FT3) upon full-text inspection.
EC5	Unrefereed preprint/deposit lacking peer-review venue metadata.

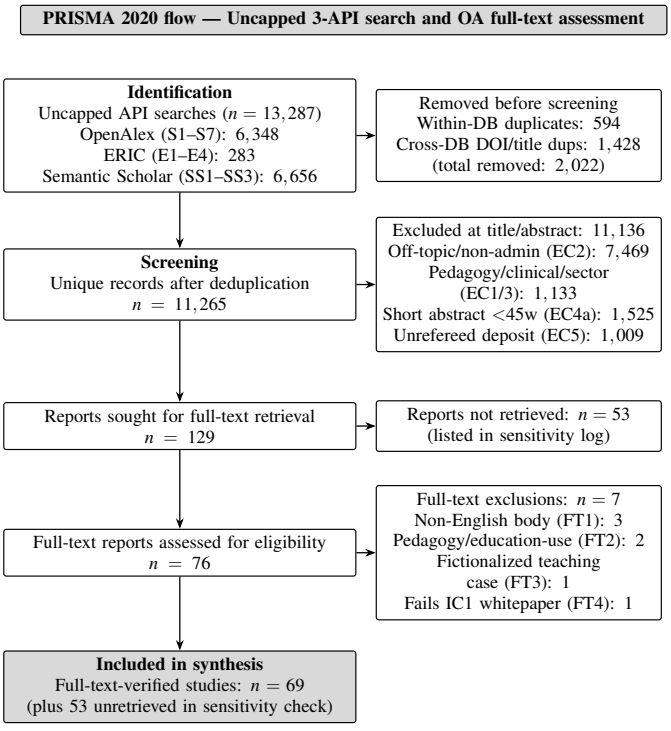


Fig. 2. PRISMA 2020 flow diagram: 13,287 records identified, 11,265 screened, 129 sought for retrieval, 53 not retrieved, 76 full texts assessed, 7 excluded, and 69 included.

category counts indicate thematic prevalence rather than effect size. Barrier rank order was also checked in a 122-study sensitivity pool comprising the 69 full texts and the 53 reports sought for retrieval but not obtained as verified English full texts.

G. Computational Evaluations of the Architecture (Not SLR Evidence)

RQ8 is not answered by the systematic review. Four computational evaluations are reported separately in Section VII and Appendix D and are not pooled: (A) RA-GO 20+30 specification-conformance runs; (B) rule-based BPI 2019 information access on the closed-case D-sample and S/S2 injections; (C) language-model orchestration on a frozen 138-record BPI test partition with independent D/S/clean-by-rule labels that are distinct from Evaluation B's D-sample; (D) 34 scenario-based modernization cases with gold locked before scoring. Methods for C and D (model, seed, hashes, gold lock, held-out faults) are specified in those subsections and the appendix.

IV. RESULTS

This section reports the review evidence that motivates the six principal findings of Table I; the findings themselves are established in Sections V-B and VIII. The corpus documents approval barriers consistently but rarely measures approval time or outcome, so F1–F3 are tested on event logs rather than synthesized from the literature alone.

A. Study Selection

The 14 uncapped searches returned 13,287 records (OpenAlex: 6,348; ERIC: 283; Semantic Scholar: 6,656). Removing 2,022 duplicates (594 within-source, 1,428 cross-database) left 11,265 unique records. One reviewer advanced 129 reports to full-text retrieval, with title/abstract exclusion reasons for the remaining 11,136 records coded as in Section III (`EC2` 7,469; `EC1/EC3` 1,133; `EC4a` 1,525; `EC5` 1,009). Of the sought reports, 76 full texts were assessed and 53 were not retrieved; of the assessed reports, 7 were subsequently excluded (three non-English bodies, two pedagogy/education-use studies, one fictionalized teaching case, and one ACODE survey whitepaper failing IC1), yielding 69 verified English full-text studies (Fig. 2).

B. Characteristics and Formal Quality Appraisal

Table V summarizes the 69 full-text-verified studies (item-level MMAT/JBI scores in Appendix C, Tables XXVII–XXVIII). Exactly 53 examine higher education and 16 examine comparable public-sector administrations across 2016–2026. Geographically, 21 are set in Asia or the Middle East, 19 are global or non-site-specific reviews/frameworks, 10 in Africa, 9 in Europe, 7 multi-region, 1 in the United States, and 2 in Latin America. Among the 53 reports sought but not retrieved as verified English full texts, several have U.S. settings on abstract [81]–[84].

TABLE V
STUDY CHARACTERISTICS AND MMAT 2018 / JBI ITEM APPRAISAL ($N = 69$)

Characteristic	Studies (n)
Sector	
Higher education (HE)	53
Public-sector administration comparator (PS)	16
Study design family (verified from article body)	
Qualitative / case / process study	18
Quantitative / survey	12
Systematic / scoping / integrative review	13
Conceptual, legal or architectural analysis	11
Mixed methods	10
Design science research	5
Geographic setting	
Asia / Middle East	21
Global / not site-specific	19
Africa	10
Latin America	2
Europe	9
Multi-region	7
United States	1
*Author-defined descriptive summary bands (% checklist items met)**	
$\geq$80% of applicable criteria met	46
60–79% of applicable criteria met	18
$<$60% of applicable criteria met	5
Appraisal instrument breakdown (mean % items met)	
MMAT 2018 Qualitative (mean 81%)	19
MMAT 2018 Quant. Descriptive/NR (mean 58%)	12
JBI Systematic Review Checklist (mean 73%)	13
JBI Text/Legal/DSR (mean 93%)	15
MMAT 2018 Mixed Methods (mean 82%)	10

*MMAT 2018 guidance [76] recommends reporting item-level responses rather than official overall quality cutoffs; the $\geq$80%/60–79%/$<$60% bands above are author-defined descriptive summaries only, with all individual item scores and verbatim passages preserved in the archive and Appendix C. Reference lists were stripped prior to body-text verification.

On item-level MMAT 2018 and JBI appraisal, 46 studies met $\geq$80% of applicable criteria, 18 met 60–79%, and 5 met $<$60% (author-defined descriptive bands; item scores in `formal_mmat_jbi_appraisal.csv`). Among the 13 systematic/scoping reviews, the most frequent deficits were absent dual-reviewer appraisal (Q6) and unassessed publication bias (Q9); among the 12 quantitative surveys, explicit nonresponse-bias testing (Item 4.4) was the primary weakness. Full texts on Namibian leadership [72], Brazilian IT-governance maturity [73], and administrative digital transformation [74] identify leadership and culture as binding constraints and do not measure approval time.

C. Governance and Decision-Making Approaches (RQ1)

Table VI synthesizes six governance approaches. Across full-text studies and abstract-level sensitivity comparators, centralized vs. decentralized authority exhibits an outcome-specific trade-off: in U.S. panel data retained only at abstract level [81], centralized IT decision rights correlate with fewer cybersecurity breaches in heterogeneous research universities, whereas full-text public-sector and university qualitative studies show that rigid, compliance-driven centralization generates procurement bottlenecks, informal workarounds, and legitimacy erosion [85]–[87]. Multi-case full-text university studies identify baseline structural, procedural, and relational IT-governance mechanisms [88], [89] and agile governance structures [90], [91], while warning that under-resourced committees often fail to operationalize formal mandates [92], [93]. A verified Namibian mixed-methods study finds dedicated leadership and collaborative networks enabling IT governance [72]; a Brazilian documentary comparison reports COBIT-adapted maturity of 76.67%, 63.33%, and 0% across three federal universities [73]. The evidence thus supports *risk-calibrated conditional governance* rather than uniform centralization or deregulation.

D. Barriers to Administrative Modernization (RQ2)

Across the 69 full texts, seven non-exclusive barrier categories were verified (Fig. 3, Table VII): governance capacity and leadership (47 studies, 68%), technological and legacy fragmentation (29, 42%), structural and administrative approval arrangements (25, 36%), cultural and human factors (21, 30%), AI and decision-support governance (22, 32%), financial and procurement constraints (12, 17%), and political/legitimacy pressures (12, 17%). Pooling the 53 abstracts of reports not retrieved in full text (122 studies total) preserves the same top-three barrier ordering (B6, B4, B1). Full-text inspection sharpens two mechanisms: first, staff resistance frequently reflects opaque, top-down rollout speed, poor ERP–process fit, and change-management deficits rather than innate technophobia [6], [7], [93], [111]; second, digitizing unexamined workflows risks "digital bureaucracy" or "automation traps" unless approval thresholds and ownership silos are redesigned [102], [112].

E. Theoretical Lenses (RQ3)

Among the 69 full texts, 20 draw on IT-governance, enterprise-architecture, or portfolio constructs [88]–[90], [113], [121], [124], [127], 11 on algorithmic, legal, or human-centered AI governance [64], [106], [115], [129]–[132], [134], [135], 4 on sociotechnical or human-centered transformation [106], [120], [136], [137], and 6 on institutional theory or coercive/normative pressures [4], [7], [116]. Several newly retrieved empirical studies combine Technology–Organization–Environment (TOE) or dynamic-capability models with administrative process or RPA/ERP adoption [117], [118], [125]. Fig. 4 synthesizes how legacy structures, distributed authority, and statutory safeguards jointly shape approval processes and workaround behavior.

F. Relation to Control Frameworks (RQ4)

Among the 69 verified article bodies, 13 explicitly apply or analyze COBIT [73], [85], [86], [88]–[92], [97], [121], [138], 10 reference ISO/IEC 38500 [113], 5 apply TOGAF enterprise architecture [86], [102], [120], [121], [138], and 1 integrates NIST alongside COBIT and ISO/IEC 38500 [90]. Fig. 5 and Table VIII map B1–B7 to COBIT 2019 objective titles [32] and verified NIST CSF 2.0, AI RMF 1.0, AI 600-1, SP 800-53 r5, and SP 800-161 r1 controls [33]–[37]. COBIT entries use published objective titles only.

TABLE VI
CRITICAL COMPARISON OF GOVERNANCE AND DECISION-MAKING APPROACHES (RQ1)

Approach	Reported benefits	Reported limits or conflicting evidence	Evidence base and indicative confidence
Centralized IT decision authority	Fewer breaches, with larger effects for heterogeneous infrastructure and public or research-intensive universities [81]; standardization and security benefits [3].	Control-oriented centralized procurement linked to inefficiency, workarounds, and declining legitimacy [94]; dependencies and lags [3]; flexible security management perceived as more effective in decentralized, shared-governance institutions [84].	One large observational US panel (moderate) [81]; other evidence qualitative or perception-based (low). Outcomes differ (security versus efficiency), so findings are not directly contradictory.
Delegated and participatory committees	Delegated decision model and inclusive committees associated with a stronger innovation culture [82]; faculty voice articulated through senate committees [95].	Committees struggled to operationalize their mandate in a low-resource setting [92]; prior evidence ambiguous on whether mature governance helps or hinders innovation [82].	Small multi-case and practice reports (low to very low).
Framework- and maturity-driven governance (COBIT, ISO/IEC 38500 style)	Baseline practice models for universities [88], [96]; IT-governance mechanisms positively related to absorptive capacity [97]; diagnostics reveal large maturity differences [73].	Maturity often low [98]; governance structures not directly associated with transformation, with cybersecurity mediating and institutional decoupling reported [99]; document-based diagnostics may not reflect practice [73].	Mostly descriptive assessments and surveys; little outcome evidence (low).
Process redesign with regulation-encoded workflows	Integrating dispersed data, encoding regulatory constraints, and short feedback cycles shortened a regulation-intensive process [100]; purchasing process modeled per regulation [101].	Proposals such as threshold-based approvals and automated budget tracking are untested [102].	One practice-based single case, one design-science study, one process analysis (low).
Procurement maturity and reform	Procurement maturity described as a catalyst for IT procurement success [103]; reform pathways include streamlined procurement and capacity building [86].	Software procurement bypassing strict practice, weak contracts [87]; joint procurement only partly alleviates difficulty [104].	Qualitative cases and interviews; model in [103] needs validation (low).
AI-assisted decision support	Possible gains in scalability, cost, and quality [105]; reported production use with shorter response preparation and preserved human oversight in one public administration [106]; efficiency and analytics gains when readiness and ethical governance exist [107].	Opacity, surveillance, and centralization of managerial authority [108]; equity, manageability, and political-feasibility concerns [105]; isolated safeguards give limited assurance [109]; performance figures in architecture proposals not independently verifiable [110].	Mostly conceptual, review, or self-reported (low to very low). No independent outcome evaluation identified.

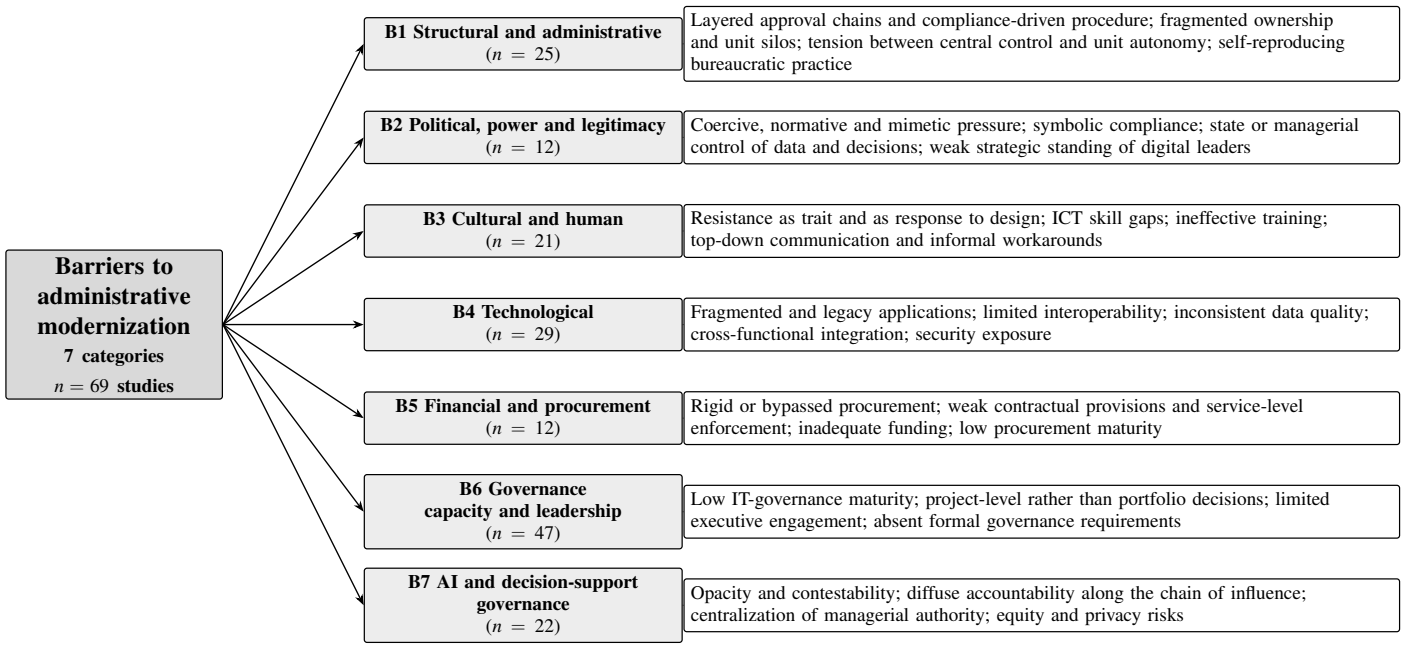


Fig. 3. Taxonomy of barrier categories B1–B7 with counts of full-text-verified included studies coded to each category ($N = 69$, non-exclusive).

TABLE VII
BARRIERS TO ADMINISTRATIVE MODERNIZATION VERIFIED FROM FULL TEXTS ($N = 69$; RQ2)

ID	Category	n	Representative full-text-verified evidence	Methodological caveat (MMAT/JBI)
B1	Structural and administrative	25	Multi-stage university approval bottlenecks across ∼100 BPMN processes [102]; strategy–operations misalignment and manual workarounds [93]; multi-step HR bureaucracy [112]; provincial compliance delays [85]; IT portfolio–committee gaps in Spanish universities [113].	Cross-sectional or single-institution process mappings; few longitudinal designs.
B2	Political, power, and legitimacy	12	Four response modes to coerced transformation including surface compliance and active resistance [7]; state centralization of university data power [114]; university autonomy vs. external AI regulation [115]; algorithmic managerialism [108]; situated legitimacy work in municipal digital transformation [116].	Primarily qualitative interview, doctrinal, and critical policy studies.
B3	Cultural and human	21	Resistance driven by speed and opacity of change [6]; intolerance of ambiguity [5]; staff skill and change-management gaps in university ERP and RPA rollouts [111], [117], [118].	Heavy reliance on self-report surveys; common nonresponse reporting gaps (MMAT 4.4).
B4	Technological (legacy, integration, data, security)	29	Legacy monolith decoupling constraints [119]; fragmented university systems and interoperability gaps [91], [120], [121]; data quality barriers [122]; cloud-ERP security gaps [123]; Initial digital-readiness bottlenecks in a Brazilian public university [124].	Architectural evaluations often lack multi-institution comparative baselines.
B5	Financial and procurement	12	Bypassed university software procurement procedures, weak SLAs, and underfunding [87]; rigid public ICT procurement [86]; beyond-acquisition organizational gaps [125]; resource scarcity in AI/EA adoption [121], [126].	Few studies link procurement rules to objective cycle-time or defect logs.
B6	Governance capacity and leadership	47	Dispersed ITG committees in low-resource HEIs [92]; project-level vs. portfolio governance [127], [128]; baseline mechanism gaps [88], [89]; agile governance capability needs [90], [97]; public AI governance capacity gaps [129]–[131].	Most frequent category (47/69), partly reflecting IT-governance search terms.
B7	AI and decision-support governance	22	Accountability along chains of influence [64], [132]; explainability and university constraint feasibility rated hardest by experts (I-CVI=0.71) [133]; policy–evaluation gap [15], [129], [131], [134]; HITL multi-agent drafting with mandatory human validation [106]; unverified agentic projections [110].	No randomized or controlled institutional trial of AI-assisted software approval.

Non-exclusive full-text coding ($N = 69$). In abstract-level sensitivity pooling across all 122 studies (adding the 53 reports sought but not retrieved in full text, e.g., [3], [81], [94]), the top categories remain B6, B4, and B1 in identical rank order.

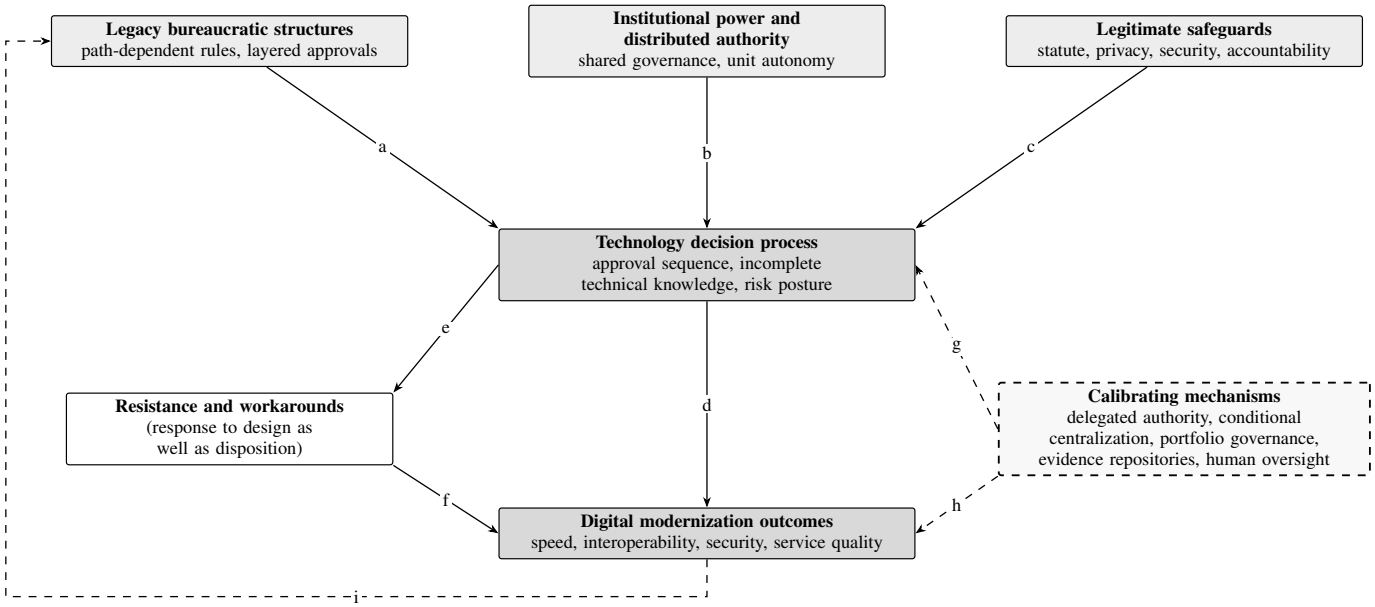
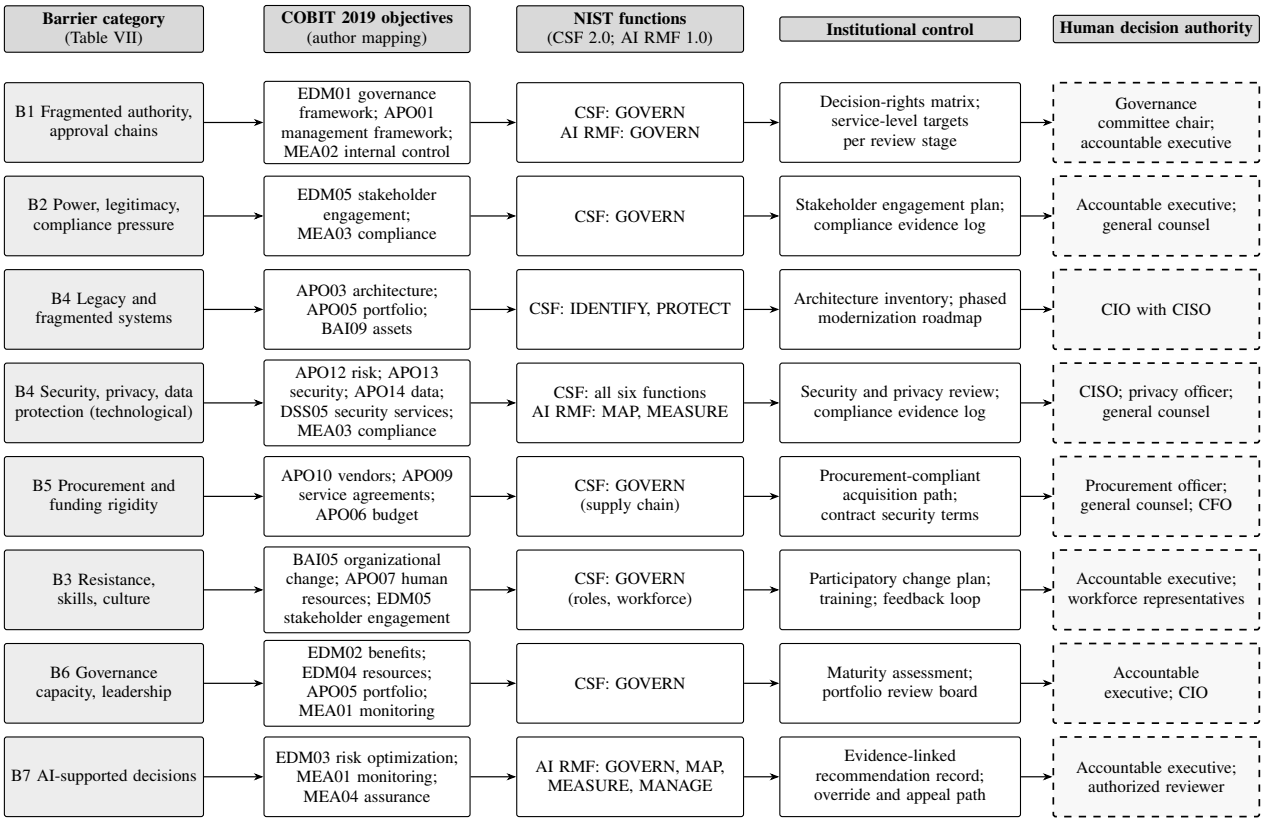


Fig. 4. Conceptual synthesis. Links: (a) legacy structures shape the decision process; (b) distributed authority shapes it; (c) safeguards constrain it; (d) the process shapes modernization outcomes; (e) the process provokes resistance and workarounds; (f) resistance affects outcomes; (g, h) calibrating mechanisms moderate the process and outcomes; (i) outcomes feed back to legacy structures.



Dashed boxes mark roles that retain legal and institutional decision authority. The crosswalk uses published objective titles and function names.

Fig. 5. Crosswalk from barrier categories to COBIT 2019 objective titles, verified NIST functions/controls, institutional controls, and human decision authorities.

TABLE VIII
ANALYTICAL CROSSWALK FROM BARRIER CATEGORIES TO COBIT 2019 OBJECTIVES AND NIST OUTCOMES (RQ4)

Barrier	COBIT 2019 objectives (published titles)	NIST outcomes (verified against source text [33], [34], [36], [37])	Mapping rationale
B1	EDM01 governance framework; APO01 management framework; MEA02 system of internal control	CSF GV.RR-02 (roles, responsibilities, and authorities established, communicated, understood, enforced); SP 800-53 PL family (planning)	Decision rights, escalation paths, and control ownership are what layered approval chains lack when authority is fragmented.
B2	EDM05 stakeholder engagement; MEA03 compliance with external requirements	CSF GV.OC-02 (stakeholders and their expectations understood), GV.OC-03 (legal, regulatory, and contractual requirements understood and managed)	Legitimacy pressures and compliance duties call for explicit stakeholder and regulatory accountability instead of symbolic compliance.
B3	BAI05 organizational change; APO07 human resources	CSF GV.RR-01 (leadership accountable; risk-aware culture), GV.RR-04 (cybersecurity in HR practices)	Resistance and skill gaps relate to leadership accountability and workforce practices.
B4	APO03 enterprise architecture; APO05 portfolio; BAI09 assets; APO13 security; APO14 data; DSS05 security services	CSF ID.AM-02 (software, services, and systems inventories); SP 800-53 CA family (assessment, authorization, and monitoring)	Legacy sprawl and security exposure relate to inventories, assessment, and authorization evidence.
B5	APO06 budget and costs; APO09 service agreements; APO10 vendors	CSF GV.SC-05 (supply-chain requirements in contracts), GV.SC-06 (due diligence before engagement), GV.SC-07 (supplier risk monitored); SP 800-53 SA and SR families; SP 800-161	Contract and vendor weaknesses relate to acquisition, supply-chain due diligence, and contract terms.
B6	EDM02 benefits delivery; EDM04 resource optimization; APO05 portfolio; MEA01 monitoring	CSF GV.RM-01 (risk-management objectives agreed by stakeholders), GV.OV-03 (performance evaluated and reviewed)	Portfolio alignment and measured outcomes address project-by-project decisions and low maturity.
B7	EDM03 risk optimization; APO12 risk; MEA01 monitoring; MEA04 assurance	AI RMF GOVERN 1.1 (legal and regulatory requirements), GOVERN 1.6 (AI inventory), MEASURE 2.8 and 2.9 (transparency and explainability); GenAI profile risks: confabulation, data privacy, information security, value chain and component integration [35]	Opaque or diffuse AI-supported decisions relate to risk ownership, inventory, explainability, and independent assurance.

COBIT entries use published objective titles only [32]. NIST entries were verified against the cited source texts.

G. Strategies (RQ5)

Table IX synthesizes ten strategy categories across the 69 full texts, led by auditable controls and standard frameworks (S8, 18 studies), human oversight and accountability by design (S7, 16), participatory and adaptive change management (S5, 13), process redesign and workflow integration (S1, 11), readiness diagnostics (S6, 11), and portfolio/EA alignment (S9, 13). Full-text evidence documents BPMN threshold redesign across ∼100 university processes [102], stakeholder-driven ERP governance across 10 German universities [111], and analytical LLM workflow redesign across 5 university HR processes [112], though none conducts a randomized controlled comparison of executive modernization-approval regimes.

H. AI-Assisted Decision Support (RQ6)

Across 22 full-text studies addressing AI governance, legal and public-administration analyses converge on mandatory human oversight, traceability along the chain of influence, transparency, and contestability for consequential decisions [115], [132], [135], [145]. Empirical and design-science university studies show that open LLMs and AI tools remain concentrated in operational automation rather than validated strategic governance [15], [134], and expert content-validity testing of an AI-EMIS architecture flagged AI explainability/traceability and feasibility under university constraints (I-CVI=0.71) as the hardest requirements to satisfy [133]. [112] evaluates Step-3 autonomous AI decision execution with analytical workflow formulas; [110], [150] report multi-agent legacy-modernization metrics as architectural projections rather than controlled institutional trials. Related work on enterprise generative-AI security and agentic architectures emphasizes oversight, benchmarks, and governance frameworks [49], [50], [61], and high-stakes financial decision-agent settings raise related cybersecurity and regulatory questions [151].

I. Practitioner Benchmarks (Gray Literature, Outside the Included Set)

Two verified practitioner sources anchor the statutory gate list and the scenario contexts used later, and neither is counted as included evidence. The May 2025 EDUCAUSE QuickPoll ($n = 270$ self-selected responses) [152] reports that 49% of institutions account for AI in procurement while only 36% re-review after procurement, citing AI pace of change (45%), insufficient institutional AI governance (40%), and opaque vendor terms (21%). The December 2025 Texas Quality Assurance Team report ($n = 59$ major state and higher-education IT projects, $1.9B) [153] documents the statutory thresholds reused in Gate H1—Project Delivery Framework monitoring above $5M, pre-execution contract review at ≥$10M, and pre-contract TX-RAMP certification—alongside a 57.6% (34/59) within-10% budget and schedule rate.

V. Discussion

A. Bureaucracy as Friction and as Safeguard

The evidence does not support treating bureaucratic structure as uniformly harmful. The same arrangements appear as obstacles to responsiveness [94], [102] and as sources of legal certainty, accountability, standardization, and security [3], [4], [81]. Institutional lenses explain why structures persist [20]–[22] but rarely distinguish protective from wasteful persistence; the included empirical studies mostly describe barriers without measuring approval time or outcome (Table X).

B. A Governance-Frontier Model of Protective and Deadweight Friction

We formalize the distinction the evidence leaves unmeasured as a mid-range model under the stated assumptions. Let $x \geq 0$ be the review intensity applied to an initiative (depth and breadth of required checks). Residual expected loss is $L(x) = L_0\,(1-\rho(x))$, where $L_0 > 0$ is the exposure without review and ρ is the share of exposure removed, with $\rho(0) = 0$, $\rho' > 0$, $\rho'' < 0$. Total delay is

$$T(x;\theta) = \underbrace{p\,x}_{\text{protective friction}} + \underbrace{(1-\theta)\,D(x)}_{\text{deadweight friction}}, \tag{1}$$

where $p > 0$ is the unavoidable time per unit of review, $D(0) = 0$, $D' > 0$, $D'' \geq 0$ is delay from queueing, handoffs, and unclear ownership, and $\theta \in [0,1]$ is coordination capacity (clear decision rights, shared evidence, parallel review). Decision makers weigh loss against delay with $\lambda > 0$ (loss-equivalent cost per unit time), so the objective is $J(x;\theta) = L(x) + \lambda\,T(x;\theta)$. The set $\{(T(x;\theta), L(x))\}_{x\geq 0}$ is the governance frontier for a given θ (Fig. 6).

Optimum. Comparing the marginal reduction in loss with the marginal delay cost at $x = 0$ shows that an interior optimum exists if and only if $L_0\rho'(0) > \lambda\big(p + (1-\theta)D'(0)\big)$; it then satisfies

$$L_0\,\rho'(x^*) = \lambda\big(p + (1-\theta)D'(x^*)\big), \tag{2}$$

and J is strictly convex because $-L_0\rho'' + \lambda(1-\theta)D'' > 0$, so x^* is unique. If $L_0\rho'(0) \leq \lambda\big(p + (1-\theta)D'(0)\big)$, then $x^* = 0$: review is not worth its cost. (The quadratic deadweight form used later in RA-GO implies $D'(0) = 0$, which recovers the special case $L_0\rho'(0) > \lambda p$.)

Proposition 1 (Dominance): For any $x > 0$, a higher θ leaves $L(x)$ unchanged and lowers $T(x;\theta)$. Reducing x lowers delay

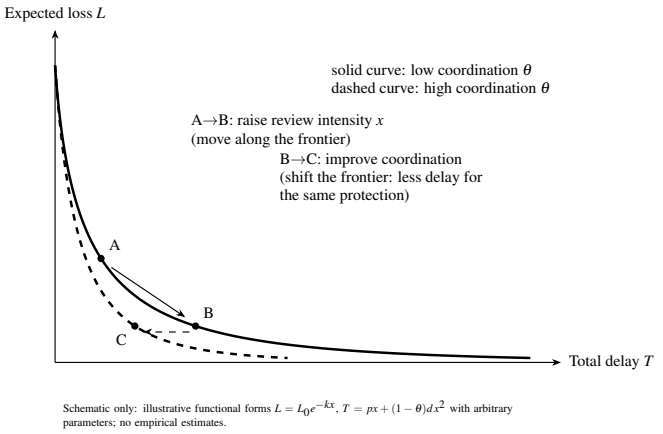


Fig. 6. Schematic governance frontier. Raising review intensity moves along a frontier; improving coordination shifts the frontier.

TABLE IX
STRATEGIES REPORTED IN FULL-TEXT STUDIES TO REDUCE DEADWEIGHT FRICTION WHILE PRESERVING SAFEGUARDS ($N = 69$; RQ5)

ID	Strategy	n	Representative full-text-verified sources	Appraisal-weighted evidence strength
S1	Process redesign and integrated workflow	11	Threshold-based approvals and BPMN automation across ∼100 university processes [102]; 5-phase AI-assisted HR workflow redesign [112]; administrative RPA and planning automation [117], [139], [140].	Observational and analytical support; no controlled trials.
S2	Governance structure design (committees, rights)	9	Baseline university structure/process/relational mechanisms [88], [89]; context-fit ITG committees [92]; agile IT governance structures [90], [91].	Design-science and multi-case qualitative support.
S3	Calibrated (conditional) centralization	2	Provincial GITO coordination vs. compliance bottlenecks [85]; public vs. private university cloud-ERP security governance [123]; U.S. panel abstract on breach reduction [81].	Context-dependent: security containment vs. local responsiveness.
S4	Procurement maturity and reform	2	Practitioner procurement-streamlining pathway [86]; university IS procurement contract/SLA enforcement [87]; related risk-aware AI procurement framing [141].	Qualitative case support; no university cycle-time log evaluation.
S5	Participatory and adaptive change management	13	Stakeholder-driven ERP rollout across 10 German HEIs [111]; digital transformational leadership and agility [142], [143]; inclusive AI governance [144], [145]; cross-country HE AI-policy readiness gaps [131].	Multi-institution survey and mixed-methods support.
S6	Readiness and maturity assessment	11	University readiness and service-planning diagnostics [118], [120], [124], [126], [133], [146].	Diagnostic reports; no causal intervention test.
S7	Human oversight and accountability by design	16	Continuous human oversight in data-centric MLOps [135]; legal contestability and transparency [64], [115], [132]; HITL multi-agent drafting with mandatory human validation [106]; ethical DDDM/AI-EMIS guardrails [133], [147].	Strong normative and legal support; limited behavioral testing.
S8	Auditable controls and standard frameworks	18	Auditable lifecycle controls with owners [148]; COBIT/ISO 38500/TOGAF/NIST integration [86], [89], [90], [113], [121]; national/public AI-initiative inventories [129], [130].	13 full texts apply COBIT; no causal test on approval speed.
S9	Portfolio and alignment governance	13	University IT project-to-governance alignment [113], [127], [128]; McFarlan portfolio mapping [146]; EA strategic alignment [121], [138].	Qualitative and descriptive survey support.
S10	Data and knowledge governance	6	Data-centric AI/MLOps governance [135]; EA knowledge integration [138]; ethical DDDM and AI-ERP data repositories [147], [149].	Design-science and single-case support.

only by raising loss, whereas raising θ lowers delay at the same loss.

Proposition 2 (Coordination): At an interior optimum, $dx^*/d\theta = \lambda D'(x^*)/\big(-L_0\rho'' + \lambda(1-\theta)D''\big) > 0$: better coordination raises the optimal review intensity, so the model predicts it can increase protection and speed together.

Proposition 3 (Comparative statics): At an interior optimum, $dx^*/dL_0 > 0$, $dx^*/d\lambda < 0$, and $dx^*/dp < 0$ (by the implicit function theorem applied to Eq. (2)). The envelope theorem gives $dJ^*/d\theta = -\lambda D(x^*) < 0$: the value of coordination equals the deadweight delay avoided.

Proposition 4 (Risk tiering): Suppose initiatives i differ in $L_{0,i}$ and the other parameters are equal. Then x_i^* is strictly increasing in $L_{0,i}$, and any uniform intensity $\bar{x}$ is suboptimal for at least one initiative whenever the $L_{0,i}$ differ. Uniform review is thus inefficient, and tiering by exposure is the efficient design.

Proposition 5 (Testable implication): If θ is observable through proxies (e.g., presence of documented decision rights or parallel review), then delay conditional on review tier contains a component that falls with θ. A process-log study finding no such component would falsify the deadweight term in Eq. (1). These propositions follow mathematically from the stated assumptions. The SLR corpus documents barriers qualitatively without estimating p, D, or θ [86], [102], so Section VIII-G tests the deadweight term of Eq. (1) on $6,886$ university payment approval cases at TU Eindhoven [154]: consolidating budget-owner and supervisor sign-off—a change in θ that leaves the required authorizations in place—is associated with -52.2% approval latency after organizational, calendar, spend, and rework fixed effects (F1), with -31.7% adjusted end-to-end process time (F2), and with null estimates in the three stages where the mechanism cannot operate (F3), while also lowering post-Admin rejections from 5.61% to 1.92%. This is the pattern Proposition 2 predicts and Proposition 5 asks for (association, not a causal effect).

C. Resistance as Response to Design

Multiple full-text studies locate resistance not only in individual ambiguity intolerance [5] but in the speed, opacity, and coercive framing of change [6], [7], poor ERP–workflow alignment [111], and misalignment between strategy and operations [93]; parallel analyses likewise treat human resistance and legacy coupling as joint modernization barriers [8], [9], [14].

D. Governance Capacity Is Often the Binding Constraint

Governance capacity and leadership was the most frequently coded barrier (47/69 full texts). Under-resourced committees [92], fragmented enterprise architecture [120], [121], project-level decisions without portfolio oversight [127], [128], and public-sector AI governance capacity gaps [106], [129], [130] indicate that layering AI onto weakly governed workflows risks automating bottlenecks rather than resolving them [112].

TABLE X
RESEARCH GAPS AND EVIDENCE STATUS AFTER FULL-TEXT SYNTHESIS AND EMPIRICAL LOG ANALYSIS (RQ7)

ID	Gap	Status in this study	Next step
G1	Controlled outcome comparison of executive modernization-approval models (cycle time, security incidents, defect rates).	Partial: UW IT portfolio ($n = 11$) [155] and Texas QAT ($n = 59$) [153] link oversight tiers to scale/risk; an unretrieved U.S. panel abstract links centralization to fewer breaches [81]; no randomized or quasi-experimental modernization-approval trial exists.	Multi-institution pre-registered comparative study with common cycle-time and incident metrics.
G2	End-to-end empirical evidence on U.S. university enterprise, AI, legacy, and infrastructure approval chains.	1 U.S.-set included full text [128]; additional U.S.-set reports among the 53 not retrieved (e.g., [81]–[84]); EDUCAUSE QuickPoll [152]; UW enterprise IT portfolio [155].	Retrieve remaining U.S.-set full texts where legally available (`reports_not_retrieved_ill_list.csv`) and collect U.S. university ticket logs.
G3	Process-mining evidence on university administrative approval workflows.	Partially closed: In TU Eindhoven ($n = 5,277$ across 36 org units) [154], consolidated BO=Supervisor sign-off is associated with -52.2% FE-adjusted post-Admin delay ($\hat{\beta} = -3.03$ d; -34.9% unadjusted / -31.7% FE-adjusted to payment request) alongside lower post-Admin rejection (1.92% vs. 5.61%). Still lacking: university cybersecurity/FERPA and broader modernization-governance event logs.	Process-mine university ServiceNow/TDX/Jira modernization and acquisition logs with explicit governance-purpose tags.
G4	Independent evaluation of AI-assisted executive decision support.	Full texts show operational automation focus [134], analytical workflow modeling [112], or unverified agentic projections [110].	Execute shadow-mode replay and prospective evaluation (`shadow_pilot_protocol.md`).
G5	Verified policy/regulatory currency mechanisms for AI advisory tools.	Untested in all 69 full texts; explainability/traceability rated hardest by experts (I-CVI=0.71) [133].	Controlled retrieval benchmarks with statutory effective-date and citation provenance checks.
G6	Human-oversight effectiveness and automation-bias mitigation in executive modernization approvals.	Normatively mandated across legal/AI full texts [115], [132]; public-sector experiments show automation bias [156].	Blinded three-arm crossover vignette experiment (`Empirical_Validation_Package/`).
G7	Stakeholder weight elicitation (w_0) for multi-criteria governance ranking.	Tested via constructed/OOD numerical experiments and Dirichlet EVPI_s triage; real executive weights not yet elicited.	Conduct Saaty AHP / swing-weighting elicitation across CIO, CISO, counsel, procurement, and stakeholder governance roles.
G8	Broader bibliographic portal coverage of this SLR.	[72]–[74] are included as verified bodies; 53 reports remain as English full texts.	

E. Implications for AI-Assisted Support

Full-text legal, technical, and empirical studies converge on designing human oversight, traceability, and contestability into administrative AI [61], [115], [132], [133], [135]. Vendor or model selection does not by itself establish legal compliance; compliance depends on data handling, contracts, and authorized human review [141]. Automated policy retrieval may reduce staleness but still requires provenance, effective dates, and human verification [57], [58].

F. Answer to RQ7: Gaps and Framework

Table X summarizes remaining research gaps, led by the absence of controlled outcome comparisons of approval and governance models.

G. How the Six Findings Answer the Central Question

Each finding in Table I addresses whether executive modernization decisions can be shortened without weakening safeguards. Evidence is staged across four approaches used here: analytical verification (F6); retrospective historical-data analysis on public event logs (F1–F3); comparative computational benchmarking of multi-expert versus Single-Agent and RA-GO arms; and scenario-based stress testing of RA-GO under misspecification and structural out-of-distribution shifts (F4–F5). Effect sizes, evidence classes, and claim boundaries are recorded in Tables I and XXIII; the payment-log rejection trade-off (5.61% vs. 1.92% post-Admin) and the assumed parameters of F4–F5 remain the principal bounds.

H. From the Findings to GAIM-DF, RA-GO, and the Optional Extension

F1–F3 locate managerial delay partly in sign-off coordination; the rejection-rate trade-off motivates non-compensatory gates; F4–F5 support tail-risk control and review triage; and F6 shows that gates, the intensity floor, and halt can bind together. Section VII reports the multi-expert benchmarks against a strong Single-Agent and RA-GO alone; a Hybrid and shadow-mode piloting remain future work.

VI. GAIM-DF: A GOVERNANCE-AWARE DECISION FRAMEWORK

GAIM-DF (Governance-Aware Institutional Modernization Decision Framework) is a human-gated advisory process for executive modernization decisions—enterprise change, legacy transformation, AI adoption, restructuring, digital infrastructure, and cross-functional governance redesign—derived from the synthesis. It specifies decision steps, hard gates, and an advisory index; this paper does not estimate institutional outcome effects of deploying the framework.

A. Design Principles

The principles follow from the evidence: (P1) treat statutory and policy safeguards as hard constraints, not negotiable preferences [4]; (P2) make decision rights and dependencies explicit before changing a step [93], [157]; (P3) calibrate centralization to risk and infrastructure heterogeneity [3], [81]; (P4) govern at portfolio level [127]; (P5) keep humans accountable and decisions contestable [109], [132]; and (P6) record evidence with provenance and effective dates so that advice can be audited. Table XI relates these mechanisms, and RA-GO v2's gates, risk-adaptive evaluation, and review triage,

to the five COBIT 2019 Evaluate, Direct and Monitor (EDM) governance objectives using published objective titles [32].

B. Ten Steps and Three Human Gates

Fig. 7 shows ten steps: identify the initiative, map stakeholders and decision rights, identify constraints, analyze friction and dependencies, assess risk, compare strategies, determine required approvals, issue an evidence-linked advisory recommendation, define indicators, and reassess on triggers. Three human gates sit at the points where authority is legal or institutional: H1 confirms hard constraints, H2 records risk acceptance by the accountable owner, and H3 is authorization by the authorized official or officials. The output of the process is an advisory record; any approval or rejection is made outside the framework.

C. Formal Definition of the Composite Index

Let $A = \{a_1, \ldots, a_m\}$ be candidate strategies. Hard gates $g_k : A \to \{0, 1\}$, $k = 1, \ldots, K$, encode legal, policy, security, privacy, accessibility, and procurement constraints, and their values are confirmed at H1. The feasible set is

$$F = \Big\{ a \in A : \prod_{k=1}^{K} g_k(a) = 1 \Big\}. \tag{3}$$

Gates are non-compensatory: no score can offset a violated gate. For $a \in F$, criterion scores $c_j(a) \in [0, 1]$, $j = 1, \ldots, J$, express normalized desirability (for risk exposure the score is $1-$ normalized exposure). Examples are expected friction reduction, risk exposure, cost, user burden, maintainability, and evidence confidence. With institution-elicited weights $w_j \geq 0$, $\sum_j w_j = 1$, the composite index is

$$S(a) = \sum_{j=1}^{J} w_j \, c_j(a), \qquad a \in F, \tag{4}$$

and is undefined for $a \notin F$. To represent uncertainty, each criterion carries an interval $[c_j^-(a), c_j^+(a)]$, giving $S^-(a)$ and $S^+(a)$ by Eq. (4); if intervals of two options overlap, the framework reports them as indistinguishable rather than ranking them. A risk veto $c_{\text{risk}}(a) \geq \tau$ with institution-set τ is a second non-compensatory condition, recorded at H2.

Weights are elicited by the organization (for example, analytic hierarchy process or swing weighting); the framework supplies no default vector. Rank stability may be assessed by drawing $\mathbf{w} \sim \text{Dirichlet}(\kappa \, \mathbf{w}_0)$ and reporting $r_{a,b} = \Pr[\text{rank}(a) = b]$. Section VIII applies this analysis only to constructed scenarios.

D. Algorithm

Algorithm 1 specifies the advisory procedure and its stopping conditions.

Algorithm 1 GAIM-DF advisory recommendation

Require: initiative I; candidate set A; gate definitions g_k; criteria c_j; institution weights $\mathbf{w}$; evidence store $\mathcal{E}$ with sources and effective dates
Ensure: advisory record R with $\text{decision} = \varnothing$
1: Map stakeholders and decision rights for I; list required approvers
2: Retrieve constraints from $\mathcal{E}$; attach source and effective date to each
3: **if** any constraint source is missing, stale, or unverifiable **then**
4: Mark R as NEEDS HUMAN REVIEW; **go to** step 12
5: Evaluate $g_k(a)$ for all $a \in A$
6: **Gate H1**: legal and policy reviewers confirm or correct g_k
7: $F \leftarrow \{a \in A : \prod_k g_k(a) = 1\}$
8: **if** $F = \varnothing$ **then**
9: Record that no feasible option exists; **go to** step 12
10: Assess friction, dependencies, and risk; apply risk veto; **Gate H2**: risk owner records acceptance
11: Compute $S^-(a), S^+(a)$ for $a \in F$; report overlapping options as indistinguishable
12: Assemble R: options, scores with uncertainty, constraints with sources, open questions, indicators, reassessment triggers; **Gate H3**: authorized official(s) decide outside this procedure
13: **return** R labeled "advisory; pending authorized decision"

E. Optional Seven-Role Decomposition

An optional decomposition uses seven advisory roles under human-controlled orchestration: intake and scope; stakeholder and decision rights; constraint retrieval; friction analysis; risk assessment; option comparison; and recommendation/audit record [49], [106], [110], [150]. Legal and security approvals remain with authorized human roles; no role writes to institutional systems.

F. Hard-Constraint Catalogue for Gate H1

Gate H1 requires a catalogue of constraints that legal and policy reviewers confirm for each initiative. Table XII lists illustrative U.S. federal and one-state sources consulted for this paper. Applicability is a legal determination. Organizations in other jurisdictions would substitute the statutes and policies that govern them; the framework treats those sources as jurisdiction-specific inputs, not as universal requirements.

G. Technical Evidence Criteria for Gates H1 and H2

For each advisory record the framework would require the following evidence, each tied to a published control vocabulary; numeric thresholds for these criteria remain to be set empirically. (i) Provenance completeness: every constraint carries its source, version, and effective date (NIST AI RMF GOVERN 1.1 on legal and regulatory requirements [34]; CSF GV.OC-03 [33]). (ii) System inventory linkage: the initiative is recorded in an inventory of software, services, and AI systems (CSF ID.AM-02; AI RMF GOVERN 1.6). (iii) Assessment evidence: security and privacy assessment and authorization artifacts exist (SP 800-53 CA, PL, and SA families [36]). (iv) Supply-chain evidence: supplier due diligence and contractual cybersecurity requirements (CSF GV.SC-05 to GV.SC-07; SP 800-161 [37]). (v) Generative-AI risk record where applicable: confabulation, data privacy, information security, and value-chain risks considered [35]. (vi) Contestability: the affected party or reviewer can override and

TABLE XI
GAIM-DF / RA-GO V2 MECHANISMS RELATIVE TO COBIT 2019 EDM OBJECTIVES (TITLES ONLY)

EDM objective (title)	Framework mechanisms (author mapping)
EDM01 Ensured Governance Framework Setting and Maintenance	GAIM-DF steps 1–3 (initiative scope, decision rights, constraint catalogue); Gate H1 confirmation of hard constraints.
EDM02 Ensured Benefits Delivery	Feasible-set comparison on benefit and outcome criteria after gates; portfolio monitoring in steps 9–10.
EDM03 Ensured Risk Optimization	Gate H2 risk acceptance; RA-GO risk caps, KL-DRO evaluation, and halt when no admissible option remains.
EDM04 Ensured Resource Optimization	RA-GO review-intensity and arrangement choice; Dirichlet EVPI_s triage of scarce human review.
EDM05 Ensured Stakeholder Engagement	Stakeholder and approver mapping; preference-uncertainty reporting; contest/override path retained by authorized humans.

Mapping uses published COBIT 2019 objective titles [32].

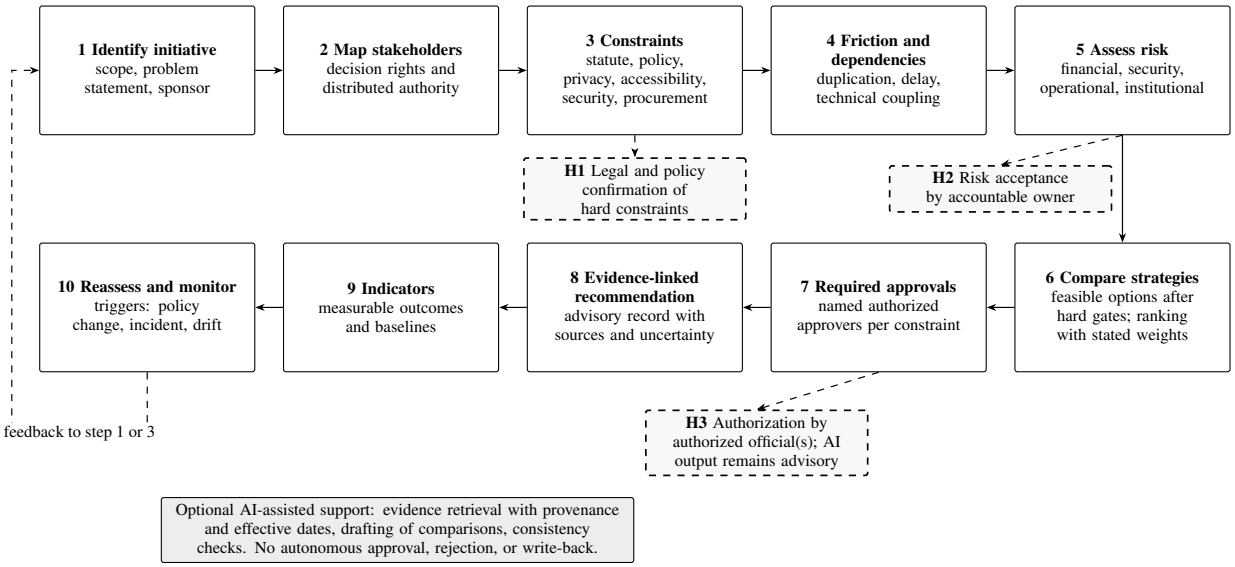


Fig. 7. GAIM-DF ten-step process with human gates H1–H3 (dashed). Optional AI support remains advisory.

appeal, with overrides logged, in the spirit of OMB's remedies and human-oversight practices [163] and GAO's governance and monitoring principles [38].

H. Mapping to Barriers and Controls

Each step addresses reported barriers: steps 1–3 address fragmented authority and constraints (B1, B2); step 4 technology and procurement dependencies (B4, B5); steps 5–7 risk and approvals (B6, B7); and steps 9–10 portfolio monitoring and learning (B6). Fig. 5 relates these to COBIT and NIST vocabularies as an interpretive crosswalk, not as a compliance determination.

TABLE XII
ILLUSTRATIVE GATE H1 CONSTRAINT SOURCES

Source	What the consulted text says (scope-limited)	Use in the framework
FERPA, 34 CFR 99.31 [158]	Prior consent is not required for disclosure to school officials with legitimate educational interests; a contractor or other outsourced party may be treated as a school official only if it performs a service the institution would otherwise use employees for, is under the institution's direct control regarding use and maintenance of education records, and is subject to the use and redisclosure requirements of §99.33(a).	Hard gate for any tool or vendor that touches education records; confirm the vendor meets all three conditions; a model or vendor choice does not itself establish compliance.
HIPAA, covered entities and business associates [159]	The Rules apply to covered entities and business associates; an entity that is neither need not comply; a covered entity must have a written business associate contract or equivalent arrangement.	Conditional gate: first determine whether the institution or unit is a covered entity or business associate.
ADA Title II web and mobile rule [160]	State and local governments, including public universities, must meet WCAG 2.1 Level AA for web content and mobile apps; after the interim final rule of April 20, 2026, compliance dates are April 26, 2027 (population 50,000 or more) and April 26, 2028 (smaller or special districts), with population determined as the rule specifies; content by contractors and third-party tools is generally covered.	Hard gate for public institutions' web and mobile software: accessibility evidence required before approval; counsel confirms the applicable date.
Procurement under Federal awards, 2 CFR 200.318 [161]	Recipients must use documented procurement procedures consistent with State, local, and tribal law, maintain oversight of contractors, avoid unnecessary or duplicative purchases, and keep records of the rationale for the procurement method, contractor selection, and price basis.	Gate for purchases charged to Federal awards: procurement record completeness and conflict-of-interest standards.
Federal Acquisition Regulation [162]	The FAR applies to all acquisitions as defined in its Part 2, except where expressly excluded.	Governs Federal agencies' acquisitions; a university's own purchases are generally governed by state law, institutional policy, and Federal-award rules, which counsel must confirm. FAR is not treated here as a direct university gate.
OMB M-25-21 [163] (replaces M-24-10 [164])	For Federal agencies, defines high-impact AI as AI whose output is a principal basis for decisions with legal, material, binding, or significant effect (including on access to education and other programs) and requires minimum practices: pre-deployment testing, a documented AI impact assessment with independent review and signed risk acceptance, ongoing monitoring, operator training, human oversight, intervention and accountability, and consistent remedies or appeals.	Benchmark, not a binding rule on universities: informs H2 evidence and the contest path for any AI-assisted advisory output.
GLBA Safeguards Rule, 16 CFR Part 314 [165]–[167]	Federal Student Aid states that institutions participating in Title IV programs are subject to the FTC Safeguards Rule through the Program Participation Agreement; revisions took effect June 9, 2023, and institutions must oversee service providers (§314.4(f)). Read via the FSA announcements and the regulation's purpose-and-scope text, not the full rule.	Conditional gate: applies where the institution handles covered customer financial information (e.g., Title IV aid); vendor oversight evidence required; counsel confirms scope.
Texas Admin. Code, 1 TAC ch. 202, subch. C [168]	Sets information security standards for Texas institutions of higher education (§§202.70–202.77), including a security-controls standards catalog based on NIST SP 800-53 Rev. 5 (§202.74) and cloud-service requirements (TX-RAMP, §202.77); the rules are being transferred to 1 TAC ch. 477. Section 202.74 text was read; the section list and transfer were taken from a secondary summary.	Applies only to institutions subject to Texas law; section numbering may change after the transfer to 1 TAC ch. 477.
Texas Responsible AI Governance Act, Bus. & Com. Code §552.001 [169]	The definition of "governmental entity" excludes institutions of higher education.	TRAIGA's governmental-entity duties do not apply to Texas public universities by that definition; counsel confirms applicability of any other state AI or vendor rules.

VII. Optional Multi-Expert Extension: Conditions, Trade-offs, and Failure Modes

This section evaluates an optional multi-expert coordination extension to GAIM-DF and RA-GO v2. The question is under what conditions, trade-offs, and failure modes the extension differs from a strong Single-Agent baseline and from RA-GO v2 alone, including under incomplete intake. Four separate evaluations are reported: (A) the RA-GO $20{+}30$ scenario comparison (Section VII-C); (B) the rule-based BPI Challenge 2019 information-access experiment, including the 89-miss and join-fix results (Section VII-D); (C) a language-model four-arm on a frozen 138-case BPI partition (Section VII-E); and (D) an independent 34-case scenario-based modernization benchmark with frozen gold (Section VII-F). Section VII-G then proposes a risk-adaptive Hybrid as a response to the observed trade-off. Evaluations A, B, and D use a deterministic rule-based prototype; C uses one 8B model on discrepancy flags. Capability decomposition and auditability requirements draw on prior multi-agent and tracing work [50], [55], [58].

A. Architecture

Executives rarely hold the legal, financial, security, procurement, operational, and AI-assurance expertise that enterprise modernization, legacy transformation, AI adoption, restructuring, digital-infrastructure, and cross-functional governance decisions need, and serialized approval chains fix the reviewers before the decision is known and bundle protective review with waiting [170], [171]. Fig. 8 shows one optional coordination design under GAIM-DF / RA-GO v2. A structured intake feeds an AI orchestrator, which convenes EXEC, GOV, and EVID always and any of seven further capabilities whose declared trigger is present (Table XIII). Each capability returns its minimum evidence, its constraint contribution, and its escalations; none approves or rejects, and none can relax a gate. GOV consolidates the contributions conservatively. With φ_k the review-intensity floor, $\kappa_k \le 1$ the risk-cap multiplier, and $\mathcal{T}_k$ the gate tokens of capability k,

$$x'_{\text{gate}} = \max\big(x_{\text{gate}}, \max_k \varphi_k\big), \qquad \text{cap}' = \text{cap} \cdot \min_k \kappa_k, \qquad \mathcal{T}' = \bigcup_k \mathcal{T}_k. \tag{5}$$

RA-GO v2 then applies the gates, optimizes under KL-ball robustness and the cap cap', filters to the Pareto set, and flags contested cases by Dirichlet EVPI. The output is one of Proceed, Pilot, Redesign, Defer, or Halt & Escalate, with a human-review flag, and the accountable executive decides within the GAIM-DF process. The ten capabilities in Table XIII are logical functions, not necessarily separate runtime agents, and should not be conflated with the six illustrative specialist roles in Fig. 1.

B. Fixing the Incomplete-Information Routing Problem

In the first prototype, a context field that was unknown at intake was treated as absent, so its trigger did not fire. This fails open: with 10% of fields unknown, a required capability was omitted in 37.4% of runs, and with 50% unknown in 96.2% (Fig. 9c). The corrected rule is fail-closed: an unknown field is treated as present, so its capability is convened, the unknown fields are listed, and the case is flagged for human review until intake is completed. The change is one line in the router, with the same activation rules otherwise. It removes the omissions (0% at every unknown-field rate tested) at a cost of 9.60–9.80 of 10 capabilities convened, so under incomplete intake the layer's routing saving over a full panel shrinks to 2–4%.

C. Comparison with RA-GO Alone and a Single-Agent Baseline

Design. Three arms use the same RA-GO v2 engine, the same 20 base and 30 structurally shifted scenarios, the same seeds, and the same Monte Carlo worlds. **R**: RA-GO v2 alone with the scenario's own gate and cap settings. **S**: a single generalist agent that applies all ten checklists to every decision, without routing. **M**: the multi-expert layer with fail-closed routing. S is a rule-based stand-in for one agent with full coverage, so the comparison isolates routing and structure rather than language-model capability. Each run is scored against the consolidated specification constraints for the scenario's full context: chosen intensity below the floor, risk above the tightened cap, or a gate-bypassing option. A second set of runs misconfigures RA-GO's own settings (cap $\times 1.5$, no intensity floor) to test where the layer should bind. Outcomes are specification-conformance proxies on constructed scenarios.

Results. On these scenarios the layer produced no measurable decision-level gain over either baseline (Table XIV). No arm violated a floor, cap, or gate, including when RA-GO's own settings were misconfigured, and the three class differences between R and M were near-tied choices that moved in both directions. S and M returned identical classes. Two differences remain. First, M convenes 0.51 fewer capabilities than a full-checklist agent under complete intake, and almost none under incomplete intake once routing is fail-closed. Second, M flags 58.2% of cases for human review against 9.5% for RA-GO alone; those flags are not validated against real faults on these scenarios. The specification-conformance check (activated set equals the trigger table's requirement in all 51 scenarios and 10,000 random vectors) confirms router consistency with the declared trigger table.

D. Does Source-Level Access Catch Discrepancies Hidden by Summaries?

As a separate information-access evaluation, we used the BPI Challenge 2019 purchase-order log [172] (CC BY 4.0; 728,558,522 bytes; MD5 `4eb909242351193a61e1c15b9c3cc814` matches 4TU) and asked: *does authorized access to source-level GR, invoice, and PO fields improve detection of discrepancies that an executive summary hides?* The log is industrial and contains no executive reports, permissions, or human discrepancy labels.

Four provenance tags are kept separate. **R**: event attributes in the XES file ($251,734$ cases). **D**: closed-case labels from van Dongen's published matching rules (3-way after/before GR, 2-way, consignment), not auditor judgments; $205,689$ closed cases in the log, of which $200,689$ enter the D-comparison

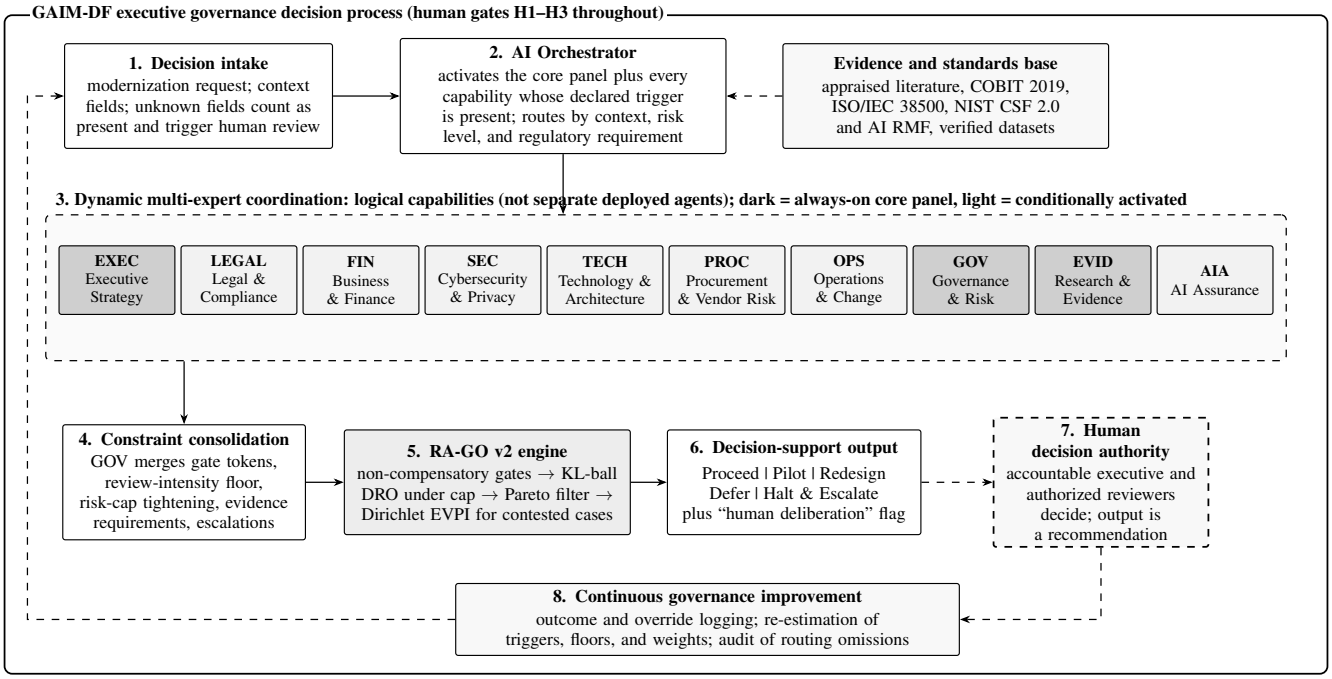


Fig. 8. Optional multi-expert coordination under GAIM-DF / RA-GO v2 with mandatory human decision authority. Unknown intake fields are treated as present (fail-closed).

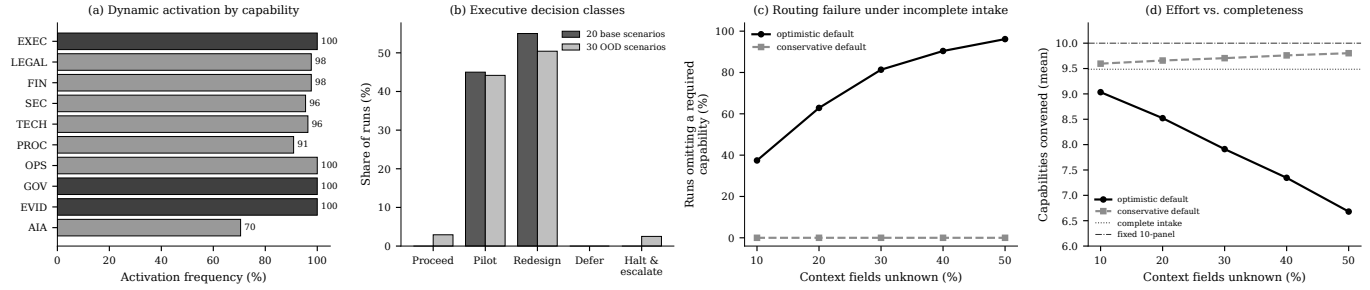


Fig. 9. Reference prototype on constructed scenarios: (a) activation frequency by capability (dark = core panel), (b) decision classes, (c) omission of a required capability when intake fields are unknown, fail-open versus fail-closed (*optimistic* and *conservative* default), and (d) capabilities convened versus the fixed ten-member panel.

after holding out the $5,000$ S-injected copies, with $10,213$ discrepant (5.1% of the D-sample), mainly missing GR ($8,205$) or missing invoice ($8,100$). **S**: $5,000$ seeded GR-drop copies; the summary is frozen from the original case and the injection list is not an input. **H**: a least-privilege matrix that forbids user identifiers, and a Halt/Proceed flag if any detector fires. The `Cumulative net worth` field is identical across events in 98.2% of cases and varies in 1.8%; additivity of invoice fragments is therefore often unobservable. 10 events carry the year 1948.

All four arms use the same documented matching checks and differ only in fields. Report-Only sees a clearance heuristic. Single-Agent sees invoices, GR, category, and user ids. Compute-matched Single-Agent and Multi-Expert each read three field groups (invoice, GR, category) without user ids; Multi-Expert partitions them across FIN, PROC, and GOV.

Source access improved detection of discrepancies the summary hid; the original Multi-Expert did *not* match a compute-matched Single-Agent on joint timestamp checks. On D, Report-Only recall is 83.6% (bootstrap 95% CI $[82.9, 84.3]$) with $1,675$ misses. Compute-matched Single-Agent recall is 100.0%. Original Multi-Expert recall is 99.1% because it missed 89 cases, all with a single D-flag (clearance before goods receipt), all in one subsidiary. FIN saw invoice timestamps, PROC saw GR timestamps, and the orchestrator only ORed flags, so the documented sequencing rule never ran. The correction is a timestamp join of four tokens already read (first invoice time, first clearance time, first GR time, and flow type): no extra source read, no added specialist, no user identifiers, and no amounts shared across FIN and PROC. Revised Multi-Expert D-recall is 100.0% (FN $= 0$) and equals the compute-matched Single-Agent. On S (GR-drop), Report-Only recall is 0.14%; original and revised Multi-Expert and both Single-Agent arms recall 100%. Only unconstrained Single-Agent read user ids (100%).

A time holdout (first event on or after 1 October 2018; $56,801$ cases, $21,705$ evaluable) was used because all 89 misses sit in one subsidiary, making a company split unusable. Two of those 89 cases fall in the holdout window. On independent S2 injections that force clearance before GR

TABLE XIII
SPECIALIST CAPABILITY REGISTER OF THE DYNAMIC MULTI-EXPERT LAYER: SCOPE, ACTIVATION TRIGGER, REQUIRED EVIDENCE, CONSTRAINT AUTHORITY, ESCALATION RULE, AND ACTIVATION FREQUENCY ON 640 CONTROLLED COMPUTATIONAL PROTOTYPE RUNS

ID	Capability	Scope	Activation trigger	Required evidence (minimum)	Constraint authority / escalation	Act. %
EXEC	Executive Strategy	Decision objectives, strategic fit, organizational value, accountability ownership	Always active (core panel)	written decision objectives and success measures; named accountable executive and decision rights basis; strategic-fit statement against the organization's plan	Proposes criterion weights; cannot relax any gate; *escalates:* Objectives conflict across business units	100.0
LEGAL	Legal and Compliance	Applicable statutes and regulations, contractual obligations, legal escalation	regulated, sq breach, uncertain applicability	citation and effective date of each asserted obligation; applicability determination by authorized counsel; contract clauses governing data, audit, and termination	Sets statutory gate tokens and a review-intensity floor; *escalates:* Applicability uncertain, or status quo itself breaches a mandate	97.7
FIN	Business and Finance	Cost, benefit, budget impact, financial uncertainty, amortization horizon	high spend	cost build-up with uncertainty range; budget source and multi-year run-rate; decision volume used for amortization (decisions per year)	Supplies cost and value distributions; no gate authority; *escalates:* Benefit case depends on unverified vendor claims	97.7
SEC	Cybersecurity and Privacy	Security exposure, privacy obligations, data protection and residency	security relevant, personal data, ai involved	data classification and flow inventory; control assessment against a named control catalogue; incident and breach-notification responsibilities	Sets review-intensity floor; may tighten the institutional risk cap; *escalates:* Regulated personal data combined with an unassessed component	95.5
TECH	Technology and Architecture	Technical feasibility, legacy compatibility, scalability, interoperability	legacy involved, high complexity	integration and dependency map; migration or wrapping feasibility assessment; exit and disposition plan for the legacy component	Supplies implementation time and complexity; no gate authority; *escalates:* No feasible integration path within the decision horizon	96.4
PROC	Procurement and Vendor Risk	Acquisition rules, supplier evaluation, third-party and supply-chain risk	vendor involved	competition or sole-source justification record; supplier due-diligence and subprocessor list; flow-down of security and privacy obligations	Sets the procurement-record gate token and a review floor; *escalates:* Vendor refuses required contractual terms	90.9
OPS	Operations and Change Management	Workflow redesign, workforce readiness, adoption, continuity	workflow change, readiness gap	as-is and to-be process description; readiness and training assessment; continuity and rollback plan	Supplies coordination capacity and readiness fit; *escalates:* Required readiness exceeds assessed capacity	100.0
GOV	Governance and Risk	Decision rights, mandatory controls, oversight, risk accountability	Always active (core panel)	mandatory-control register with owners; risk appetite and cap applicable to this decision class; oversight and monitoring plan after approval	Consolidates all floors and caps; owns the non-compensatory gate set; *escalates:* No admissible option satisfies the cap (mandatory halt)	100.0
EVID	Research and Evidence	Source verification, evidence synthesis, comparison of alternatives	Always active (core panel)	verifiable source for each material claim; appraised strength of evidence per criterion; documented alternatives considered and rejected	Flags criteria lacking verifiable evidence; *escalates:* A decisive criterion rests on unverifiable evidence	100.0
AIA	AI Assurance	Model limitations, bias, explainability, reliability, human oversight	ai involved	intended use, out-of-scope use, and known failure modes; evaluation results on in-domain data with uncertainty; logging, contestability, and override design	Sets an AI review floor; may tighten the risk cap; *escalates:* High-impact AI use affecting individual rights or entitlements	70.5

Core-panel capabilities (EXEC, GOV, EVID) are always convened so that objectives, decision rights, and evidence verification are never skipped. Only GOV consolidates review-intensity floors and risk-cap tightening into the non-compensatory constraint set; no capability can relax a mandatory gate, and none approves or rejects a decision.

($4,000$ injected + $4,000$ clean, seed 20261116), original Multi-Expert recall is 0%, while revised Multi-Expert and the compute-matched Single-Agent recall 100% with 0 false positives; Report-Only recall is 0.10%. GR-drop detection on the holdout is 100% for all source-access arms. These are information-access results on rule-derived and injected labels and are reported separately from Evaluation C.

E. Supplementary Language-Model Evaluation on BPI Challenge 2019

Evaluation C asks whether four language-model *orchestration* arms differ in discrepancy *flag* quality on a frozen 138-record test partition of the same log, with the same authorized record view. Arms are labeled LLM-RO, LLM-SA, LLM-CM, and LLM-ME. The shared output is a recommendation to route a record to human review, not an approval. Labels are 60 clean-by-rule records, 30 real cases with a documented-rule (D) violation, and 48 real records with one seeded injection (S). D labels here are distinct from Evaluation B's D-sample and

TABLE XIV
MULTI-EXPERT LAYER (M) VERSUS RA-GO ALONE (R) AND A SINGLE FULL-CHECKLIST AGENT (S) ON THE EXISTING CONSTRUCTED SCENARIOS

Outcome	R	S	M	Reading
*Calibrated settings, complete intake (*440 *runs)*				
Floor, cap, or gate violations	0	0	0	No arm violated a constraint
Decision class same as M	99.3%	100%	—	R differed from M in 3 runs, in near-tied options, in both directions
Capabilities convened (of 10)	0	10	9.49	M saves 0.51 capability vs. S (base 1.05, 95% CI $[0.65, 1.45]$; shifted 0.07, CI $[0.00, 0.17]$)
Human-review flag rate	9.5%	58.2%	58.2%	Higher flag rate under M; not validated against real faults
*RA-GO settings misconfigured (*240 *runs)*				
Violations	0	0	0	The layer's constraints did not bind either
Decision class same as M	98.8%	100%	—	As above
Human-review flag rate	8.3%	57.5%	57.5%	As above
30% *of intake fields unknown (*150 *runs)*				
Required capability omitted	—	0%	fail-open 81.3%; fail-closed 0%	Fail-closed removes the omission
Capabilities convened	—	10	8.09; 9.75	Fail-closed convenes 2.5% fewer than S
Decision class same as complete intake	—	—	100.0%; 99.3%	Omissions did not change the class either
Human-review flag rate	—	—	44.0%; 100.0%	Fail-closed sends every incomplete case to a human

TABLE XV
INFORMATION-ACCESS EVALUATION ON BPI CHALLENGE 2019. **D**: RULE-DERIVED MATCHING LABELS. **S**: GR-DROP INJECTIONS WITH FROZEN SUMMARIES. **S2**: CLEAR-BEFORE-GR INJECTIONS ON THE TIME HOLDOUT. D, S, AND S2 REPORTED SEPARATELY.

Arm	D recall (FN)	S recall	S2 holdout recall	Reads	User ids (H)	Reading
Report-Only	83.6% (1,675)	0.14%	0.10%	1	0%	Clearance heuristic; hides GR gaps and sequencing.
Single-Agent	100.0% (0)	100%	—	4	100%	Matching rule on all groups, including forbidden user ids. S2 not re-scored.
Compute-matched SA	100.0% (0)	100%	100%	3	0%	Three groups, no user ids.
Multi-Expert (original)	99.1% (89)	100.0%	0%	3	0%	Flag-OR only: misses joint timestamp checks.
Multi-Expert (revised)	100.0% (0)	100%	100%	3	0%	Same three reads; orchestrator joins FIN/PROC time tokens.

Closed cases in the log: 205,689. D-comparison sample: 200,689 (the 5,000 S-injected copies are held out of D). Discrepant: 10,213 in both frames. S: 10,000 cases in the original run. Time holdout (first event $\geq$ 2018-10-01): 56,801 cases, 21,705 evaluable, 2,002 D-positive (original ME FN = 2). S2: 4,000 injected + 4,000 clean, seed 20261116. MD5 `4eb909242351193a61e1c15b9c3cc814`.

are not human-audited anomaly labels. Frozen hashes, per-type counts, and confusion matrices (Fig. 10) are in Appendix D.

LLM-ME raised recall (93.6% vs. 70.5% for LLM-SA) at the cost of a 70.0% false-positive rate ($42/60$ clean-by-rule cases) versus 21.7% for LLM-SA, lower overall accuracy (65.9% vs. 73.9%), and about $3.9\times$ the tokens and $3.0\times$ the sequential latency (4,828 tokens / 41.9 s vs. 1,228 / 14.1 s). F1 scores are essentially tied (0.756 vs. 0.753). The case-level correctness comparison is not statistically significant (exact McNemar $p = 0.16$; Holm-adjusted $p = 0.64$). Recall on real D was $30/30$ for LLM-SA and LLM-CM and $29/30$ for LLM-ME. Extra LLM-ME detections were concentrated on injected types, including types not covered by the documented rules (I3/I5/I6: $21/24$ vs. $8/24$). LLM-RO flagged 137 of 138 cases (FPR 98.3%). LLM-CM matched LLM-SA on the flag decision in 137 of 138 cases (McNemar $p = 1$) at 85% of LLM-ME's tokens. No language-model arm beat a circular rule-engine reference on false positives ($0/60$). Type attribution was weak for every arm (58–67% type-correct recall).

F. Independent Scenario-Based Modernization Benchmark

Evaluation D compares Multi-Expert orchestration with a strong Single-Agent on *decision class* and escalation when gold labels and evidence requirements are frozen *before* either arm is run, on 34 constructed executive modernization cases (legacy transformation, digital infrastructure, AI adoption, procurement, cyber, legal, finance, incomplete departmental files; sparse and dense intake). Gold is independent of Evaluation A's specification-conformance scoring. Halt faults used here are held out from the production GOV registry of Evaluation B's join predicates. Both S and $\mathrm{M_{join}}$ apply the same operational detector; S reads the authorized token union; $\mathrm{M_{join}}$ concatenates tokens of intake-activated specialists; $\mathrm{M_{or}}$ is a no-join diagnostic. Unauthorized identifiers were present in a temptation payload.

S and $\mathrm{M_{join}}$ each classified $34/34$ cases correctly (100% accuracy; paired $\Delta = 0$). $\mathrm{M_{or}}$ classified $26/34$ correctly (76.5%) and missed 8 held-out cross-specialist HALT cases. Escalation precision and recall were 1.0 with FPR 0 for S and $\mathrm{M_{join}}$ against the frozen human-review labels (16 positives, 18

TABLE XVI
EVALUATION C: LANGUAGE-MODEL FOUR-ARM COMPARISON ON A FROZEN BPI CHALLENGE 2019 TEST PARTITION (138 CASES; ARMS LLM-RO, LLM-SA, LLM-CM, LLM-ME). LABELS: 60 CLEAN-BY-RULE, 30 RULE-DERIVED D POSITIVES, 48 SEEDED INJECTIONS (S). DISTINCT FROM EVALUATION B'S D-SAMPLE. WILSON 95% CIS IN APPENDIX D.

Arm	Precision	Recall	FPR (60 clean)	F1	Accuracy	Reading
LLM-RO	56.9%	100%	98.3%	0.726	57.2%	Flags 137/138; recall is not discriminative.
LLM-SA	80.9%	70.5%	21.7%	0.753	73.9%	Lowest FPR among LLM arms; misses most non-rule injections.
LLM-CM	82.1%	70.5%	20.0%	0.759	74.6%	Token-matched to LLM-ME; McNemar vs. LLM-SA $p = 1$.
LLM-ME	63.5%	93.6%	**70.0%** (42/60)	0.756	65.9%	Higher recall, more false alarms, lower accuracy, higher cost.

Overall correctness LLM-SA vs. LLM-ME: McNemar exact $p = 0.16$, Holm 0.64 (not significant). Mean tokens/latency: LLM-SA 1,228 / 14.1 s; LLM-ME 4,828 / 41.9 s. No LLM arm was superior on overall correctness; the comparison is a discrepancy-flagging benchmark, not an executive field evaluation.

LLM-RO ($n = 138$)

Actual class	Predicted class: Positive	Predicted class: Negative
Positive	TP **78**	FN **0**
Negative	FP **59**	TN **1**

LLM-SA ($n = 138$)

Actual class	Predicted class: Positive	Predicted class: Negative
Positive	TP **55**	FN **23**
Negative	FP **13**	TN **47**

LLM-CM ($n = 138$)

Actual class	Predicted class: Positive	Predicted class: Negative
Positive	TP **55**	FN **23**
Negative	FP **12**	TN **48**

LLM-ME ($n = 138$)

Actual class	Predicted class: Positive	Predicted class: Negative
Positive	TP **73**	FN **5**
Negative	FP **42**	TN **18**

Fig. 10. Evaluation C confusion matrices on the frozen 138-case partition (raw final-call JSON). Rows are the actual class and columns the predicted class. Actual positive is a discrepancy and actual negative is clean-by-rule; predicted positive is a human-review flag. Counts match Table XVI: LLM-RO 78/59/0/1, LLM-SA 55/13/23/47, LLM-CM 55/12/23/48, LLM-ME 73/42/5/18 (TP/FP/FN/TN).

TABLE XVII
EVALUATION D: INDEPENDENT 34-CASE SCENARIO-BASED MODERNIZATION BENCHMARK (DROPOUT 0, SEED 0; GOLD LOCK PREFIX 861EB45130705752)

Arm	Correct	Accuracy	Halt FN	Reading
Strong Single-Agent (S)	34/34	100%	0	Same operational detector on the authorized token union.
Multi-Expert with join (M_{join})	34/34	100%	0	Equivalent to S under this protocol; paired $\Delta = 0$.
Multi-Expert without join (M_{or})	26/34	76.5%	8	Misses all eight held-out cross-specialist HALT joints.

Gold classes: 10 HALT, 6 DEFER, 4 REDESIGN, 5 PILOT, 9 PROCEED; 16 labeled human-review positives and 18 negatives. Escalation precision/recall for S and M_{join} equal 1 with FPR 0 because human-review gold is a deterministic function of the same decision tables. Under this protocol, Multi-Expert matched the strong Single-Agent with a join and was inferior without a join.

negatives); access-policy violations were 0. Under this protocol, gold is the shared detector on complete facts and freeze-time fairness required gold-fault owners to be intake-activated; S and M_{join} matched, and the join was required to recover the held-out HALTs.

G. Proposed Risk-Adaptive Hybrid Orchestration for Executive Decision Support

Across Evaluations A–D, a strong Single-Agent matched Multi-Expert on decision class when the same rules and a join were used (A, D) and had a lower false-positive rate and cost on the language-model flag task (C). Multi-Expert was more sensitive to some injected discrepancy types (C) and failed when a join was omitted (B, D). The Hybrid that follows is a design response to that trade-off.

The proposed workflow has six stages (Fig. 12). **Stage 1 — Initial Single-Agent screening.** A strong Single-Agent performs authorized preliminary analysis, evidence-completeness checks, and risk triage on the same least-privilege stores used by the other arms. **Stage 2 — Risk-adaptive specialist activation.** Multi-Expert capabilities are convened only under predeclared indicators: material risk, missing required evidence, cross-domain dependencies, uncertain applicability, statutory or procurement triggers, or specialist disagreement. **Stage 3 — Evidence reconciliation.** The orchestrator evaluates structured GOV predicates, accepts only declared EVID source

S ($n = 34$)

Actual	Predicted class				
	H	Df	Rd	Pl	Pr
H	**10**	0	0	0	0
Df	0	**6**	0	0	0
Rd	0	0	**4**	0	0
Pl	0	0	0	**5**	0
Pr	0	0	0	0	**9**

$\mathbf{M}_{\text{join}}$ ($n = 34$)

Actual	Predicted class				
	H	Df	Rd	Pl	Pr
H	**10**	0	0	0	0
Df	0	**6**	0	0	0
Rd	0	0	**4**	0	0
Pl	0	0	0	**5**	0
Pr	0	0	0	0	**9**

$\mathbf{M}_{\text{or}}$ ($n = 34$)

Actual	Predicted class				
	H	Df	Rd	Pl	Pr
H	**2**	0	0	**1**	**7**
Df	0	**6**	0	0	0
Rd	0	0	**4**	0	0
Pl	0	0	0	**5**	0
Pr	0	0	0	0	**9**

H = HALT, Df = DEFER, Rd = REDESIGN, Pl = PILOT, Pr = PROCEED.

S, human review

Actual class	Predicted class	
	Positive	Negative
Positive	TP **16**	FN **0**
Negative	FP **0**	TN **18**

$\mathbf{M}_{\text{join}}$, human review

Actual class	Predicted class	
	Positive	Negative
Positive	TP **16**	FN **0**
Negative	FP **0**	TN **18**

$\mathbf{M}_{\text{or}}$, human review

Actual class	Predicted class	
	Positive	Negative
Positive	TP **8**	FN **8**
Negative	FP **0**	TN **18**

Fig. 11. Evaluation D on the frozen 34-case gold lock (dropout 0). Top: five-class decision, actual class by row and predicted class by column (H=HALT, Df=DEFER, Rd=REDESIGN, Pl=PILOT, Pr=PROCEED). S and $\mathrm{M}_{\mathrm{join}}$ are diagonal $(10, 6, 4, 5, 9)$; M_{or} sends 7 true HALTs to PROCEED and 1 to PILOT (2 HALTs remain correct). Bottom: human-review gold versus predicted escalation (actual positive is the human-review label; predicted positive is escalation). S and $\mathrm{M}_{\mathrm{join}}$: TP 16, FN 0, FP 0, TN 18. M_{or}: TP 8, FN 8, FP 0, TN 18.

classes, reports conflicts explicitly, and keeps least-privilege partitions; necessary cross-specialist relationships are joined without extra field-group reads. **Stage 4 — Independent verification.** Predefined governance gates, source validation, and calibrated anomaly-verification rules are applied to separate verified concerns from unsupported warnings. **Stage 5 — RA-GO v2 evaluation.** Risk-adaptive assessment proceeds with the mandatory intensity floor, risk cap, gate tokens, and HALT option unchanged. **Stage 6 — Human decision.** The accountable executive receives the evidence, disagreements, uncertainty, escalation rationale, and recommended alternatives. AI must never replace formal approval authority.

The design targets fewer unnecessary specialist calls, lower coordination overhead, retained cross-domain sensitivity, and traceability without bypassing institutional authority. Table XVIII labels the Hybrid **Not yet evaluated**.

H. Prospective Evaluation Protocol for the Hybrid

A prospective comparison of (i) strong Single-Agent, (ii) Multi-Expert, and (iii) the Hybrid should be pre-registered before prompt or rule tuning. The three arms would receive identical authorized information, the same underlying model family, the same tools, the same governance policies, and a controlled compute budget, with outcome labels, escalation criteria, error costs (including asymmetric cost of missed HALT), and test splits frozen first. Metrics would include precision, recall, F1, FPR, accuracy; missed critical risks; decision quality where independent ground truth exists; human-review workload; governance-compliance violations; unauthorized reads; evidence traceability; latency and token cost.

I. Worked Example (Illustrative)

An organization considers AI-assisted processing of regulated personal records across sites as part of a broader enterprise modernization program; the intake marks it regulated, personal-data-bearing, vendor-involved, AI-involved, high-spend, and complex, with a readiness gap and uncertain legal applicability. All ten capabilities are convened; the statutory-authorization and procurement-record tokens are set; the intensity floor rises to 0.60 and the cap falls from 4.59 M$ to 3.90 M$; and three escalations are raised (high-impact AI, regulated data with AI, uncertain applicability). Of five alternatives, a fast-track acquisition path bypasses a statutory gate and is excluded, and the engine recommends a shadow-mode pilot at intensity 0.90 with a human-review flag. RA-GO v2 alone returns the same class; what the layer adds is the named evidence list, the tightened cap, and the escalations the executive sees first. This example is illustrative.

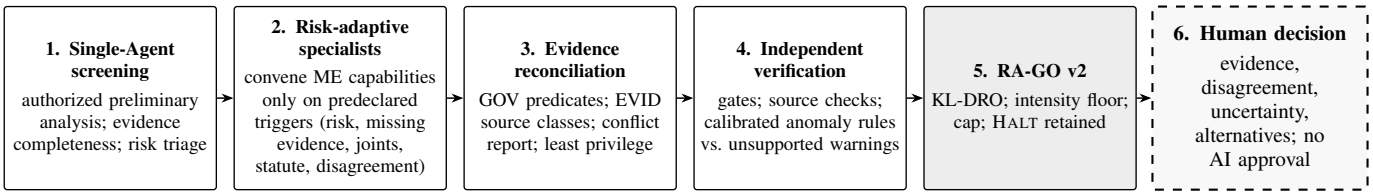
J. Summary of Multi-Expert Evaluations

Across Evaluations A–D, the optional multi-expert extension matched a strong Single-Agent and RA-GO v2 on decision class when joins and shared detectors were used, and it did not outperform either baseline overall. Incomplete-intake routing and OR-only aggregation produced identifiable failures. LLM-ME's higher recall came with a 70.0% false-positive rate, higher token cost and latency, and no significant overall correctness gain. Rule-based runs produced structured evidence lists and auditable escalations, with a modest routing saving that fail-closed intake largely removes.

VIII. RA-GO v2: Formulation and Evidence for the Principal Findings

A. Purpose and Epistemic Status

This section extends GAIM-DF with Risk-Adaptive AI Governance Optimization (RA-GO v2), which selects among retain, redesign, automate, or do-not-proceed actions while mandatory gates remain fixed, and reports evidence for the six findings in Table I. Results are reported in three classes:



Proposed Risk-Adaptive Hybrid Orchestration (design only; not evaluated). Stages 1–5 recommend; Stage 6 decides.

Fig. 12. Proposed six-stage Risk-Adaptive Hybrid Orchestration.

TABLE XVIII
COMPACT COMPARISON OF EVALUATED ARMS AND THE PROPOSED HYBRID

Property	Strong Single-Agent	Always-on Multi-Expert	Proposed Hybrid
Status in this paper	Evaluated (A, C, D)	Evaluated (A, B, C, D)	**Not yet evaluated**
Decision class vs. S (A, D)	Reference	Equal given join and the same rules	Unknown
LLM detection (C)	FPR 21.7%; acc. 73.9%	FPR 70.0%; acc. 65.9%	Unknown
Specialist activation	Authorized union / one pass	Routed or always-on panel	Risk-adaptive (Stage 2)
Cross-specialist joints	Visible if the union is read	Require an explicit join (B, D)	GOV predicates (untested)
Human authority	Recommendation only	Recommendation only	Recommendation only

analytic (F6), scenario-based stress testing on 20 base and 30 out-of-distribution (OOD) scenarios with assumed parameters (F4, F5), and observational associations on public university event logs (F1–F3).

B. Formulation

Decision sets. For scenario s, let $\mathcal{A}_s$ be the alternatives, $g \in \{0,1\}$ the governance arrangement (0 as-is, 1 redesigned to improve coordination), and $x \in [x_s^{gate}, 1]$ the review intensity. An alternative a carries a set m_a of mandatory gates it bypasses. The *gate-feasible set* is $\mathcal{F}_s = \{a \in \mathcal{A}_s : m_a = \emptyset\}$ and $x \geq x_s^{gate}$ enforces the minimum review that the gates require. Gates are hard, non-compensatory constraints: no weight, cost, or speed can trade them away.

Outcomes. With uncertain inputs ω (value V_a, cost C_a, implementation time τ_a, exposure $L_{0,a}$, coordination θ, redesign gain ξ), residual loss, delay, and cost are

$$L_a(x;\omega) = L_{0,a}\big(1-\rho_a(x)\big), \quad \rho_a(x) = \bar{\rho}\,(1-e^{-k_a x}), \tag{6}$$

$$T_{a,g}(x;\omega) = \tau_a + h_s + p_s\,x + (1-\theta_g)\,D_s\,x^2, \tag{7}$$
$$C_{a,g}(\omega) = C_a + g\,c_s^{red}, \quad \theta_g = \min\{1, \theta + 0.35\,g\,\xi\}.$$

Here $\bar{\rho} = 0.9$, h_s is the irreducible time of the $h{+}1$ mandatory human sign-offs (it is never reduced by g, which enforces the human-authorization constraint), $p_s x$ is protective friction, and $(1-\theta_g)D_s x^2$ is deadweight friction as in Section V-B. Status-quo alternatives take $x = 0$ and unreviewed loss and delay.

Robust evaluation. Let $\hat{P}$ be the empirical distribution of N joint draws. For radius $\varepsilon \geq 0$ the worst-case expectation over the Kullback–Leibler ball is

$$\rho_\varepsilon[f] = \sup_{Q:\,\mathrm{KL}(Q\|\hat{P})\leq\varepsilon} \mathbb{E}_Q[f] = \inf_{\alpha>0}\ \alpha\varepsilon + \alpha \log \mathbb{E}_{\hat{P}}\big[e^{f/\alpha}\big], \tag{8}$$

the dual form of Hu and Hong [173], used here in the framework of distributionally robust optimization [174]–[176]. We set $\varepsilon = \chi^2_{1,0.95}/(2n_{\text{eff}}) = 0.19$ with $n_{\text{eff}} = 10$ units of elicited evidence. This calibration rule is a convention and n_{eff} is an assumption; Section VIII-E varies ε.

Option selection and three-way feasibility partition. To prevent contradictions among legal gates, minimum review intensity $x_s^{gate} \in [0,1]$, and no-action options, the formulation partitions decisions into three mutually exclusive classes: (i) *active modernization options* (a, g, x) with $a \in \mathcal{F}_s^{\text{act}} = \{a \in \mathcal{A}_s \setminus \{a_{\text{sq}}\} : m_a = \emptyset\}$ and $x \in [x_s^{gate}, 1]$; (ii) *status-quo continuation* $(a_{\text{sq}}, 0, 0)$, admissible *if and only if* continuing the legacy arrangement breaches no mandatory requirement ($m_{a_{\text{sq}}} = \emptyset$) and satisfies the institutional risk cap ($\rho_\varepsilon[L_{a_{\text{sq}}}] \leq \kappa_s$); and (iii) *mandatory halt and escalation* (HALT), returned whenever the admissible cap-satisfying set is empty. Unlike our initial implementation (v1), which reverted to a_{sq} when κ_s excluded all active options even if a_{sq} itself exceeded κ_s, RA-GO v2 never silently selects a cap-violating status quo. For each $a \in \mathcal{F}_s^{\text{act}}$ and g, RA-GO chooses $x^\star = \arg\min_{x\in[x_s^{gate},1]} \rho_\varepsilon[L_a + \lambda_s T_{a,g}]$, filters all surviving active and status-quo candidates by $\rho_\varepsilon[L] \leq \kappa_s$ (returning HALT if none survive), and builds the criteria vector $z = (V, -L, -T, -C,\ \text{readiness fit},\ -\text{complexity})$ using component-wise robust evaluations (noting subadditivity $\rho_\varepsilon[L+\lambda_s T] \leq \rho_\varepsilon[L] + \lambda_s \rho_\varepsilon[T]$). After Pareto filtering, it draws $w \sim \mathrm{Dir}(c\,w_0)$ to estimate rank acceptabilities $b_o^r = \Pr(\text{option } o \text{ has rank } r)$ [177], [178] and the expected value of preference information $\mathrm{EVPI}_s = \mathbb{E}_w[\max_o w^\top \tilde{z}_o - w^\top \tilde{z}_{o^\star}]$, recommending $o^\star = \arg\max_o b_o^1$ or escalating high-EVPI_s / low-$b_{o^\star}^1$ cases to human deliberation at Gate H3. Every output remains advisory. Fig. 13 summarizes the five-layer architecture that realizes this formulation while keeping human authority outside the optimizer.

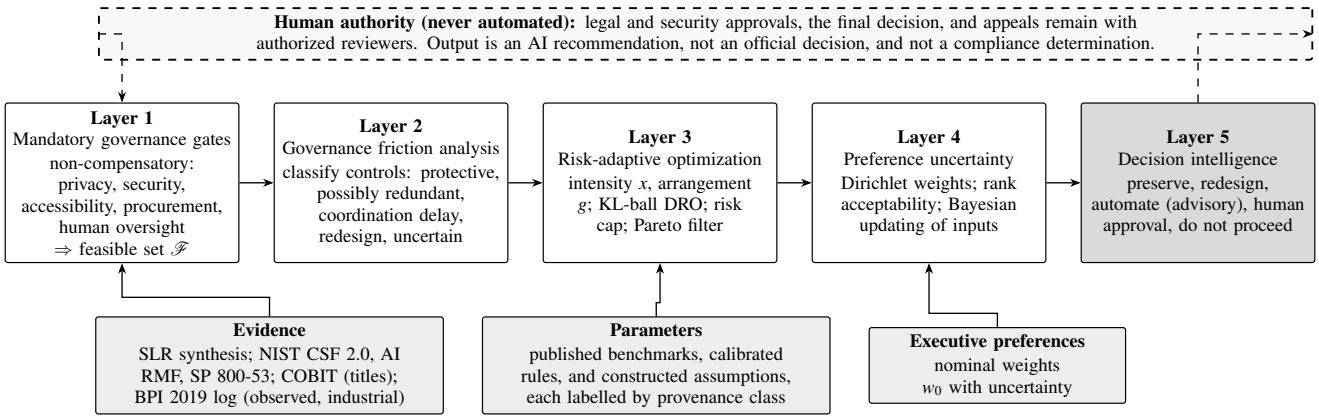


Fig. 13. RA-GO v2 integrated architecture: five layers extend GAIM-DF; human authority stays outside the optimization.

C. Scenario Design

The 20 scenarios (Appendix E, Table XXIX) span executive modernization contexts: AI adoption and acquisition (A1–A5), legacy-system transformation (B1–B5), cybersecurity, privacy, and regulatory governance (C1–C5), and administrative and cross-functional bureaucracy (D1–D5). Each has five alternatives generated from a category template and scaled by a scenario value scale, risk factor, readiness, coordination, and regulatory heaviness, with seeded jitter on template multipliers. Every scenario includes one alternative that bypasses a mandatory gate (for example an unreviewed fast-track acquisition path, a patch that skips assessment, or vendor self-attestation only), one status-quo alternative, and an AI-specific uncertainty factor where AI is involved. Scenario gate examples cite Texas or federal rules as illustrative sources only. Qualitative anchors come from the review, the EDUCAUSE QuickPoll, GAO, NIST, and legal sources, plus the BPI Challenge 2019 purchase-order log [172] (MD5-verified; 251,734 cases, 1,595,923 events), whose rework and handoff burdens—8.9% of cases repeating an activity and 73% of per-case inter-event time falling between different resources—motivate the coordination ranges without calibrating them; that log records company purchase orders, so none of its waiting times is treated as university deadweight friction outside Evaluation B. COBIT enters only through objective titles in the crosswalk (Table VIII). All modeled costs, exposures, delays, readiness values, and weights are constructed assumptions (Table XIX); F4–F6 report model behavior, while F1–F3 use separate public university event logs.

We compare RA-GO (M5) against conventional fixed-rule approval (M1: uniform intensity 0.8, gates checked, cheapest option within the risk cap), weighted-score MCDA without gates (M2), GAIM-DF (M3: gates, nominal weights, uniform intensity), GAIM-DF with a mean-risk cap (M4), an action-space-equalized GAIM-DF (M3+, defined in Section VIII-E), and the three layer ablations used in the Shapley decomposition. All methods share the same $N = 1000$ draws per scenario and seed and are evaluated on separate draws, either of a model-consistent world, of a stress world in which inputs are misspecified (heavier tails, higher exposure, weaker coordination), or of the structural OOD worlds of Section VIII-E. The stress and OOD worlds test robustness to model error, not any real hazard.

TABLE XIX
ALGORITHM AND NUMERICAL-EXPERIMENT PARAMETERS (ALL CONSTRUCTED OR CALIBRATION CHOICES)

Item	Setting
Scenarios / alternatives	20 / 5 each; arrangement $g \in \{0, 1\}$; intensity grid $x \in \{0, 0.1, \ldots, 1\}$ with $x \geq x_{gate}$
Uncertainty draws	$N = 1000$ joint draws for decisions; 3000 independent draws per evaluated world; 20 seeds (1000–1019)
Uncertain inputs	Lognormal V (CV 0.30; AI options 0.45 with Beta(7,3) performance factor), C (CV $0.25{+}0.3\,cx$), T_0 (CV 0.25), L_0 (σ=0.6); θ ∼Beta (concentration 20); redesign gain $0.35\,\xi$, ξ ∼Beta (mean = readiness, concentration 6)
Review efficacy	$\rho(x) = 0.9\,(1 - e^{-kx})$, k by alternative
DRO	KL ball around the empirical distribution; dual form; $\varepsilon = 0.19$; each metric worst-cased separately
Dirichlet MC	$w \sim \mathrm{Dir}(20\,w_0)$, $w_0 = (.25, .25, .15, .15, .10, .10)$; 2000 draws
Risk cap	$0.3\,S\,r_f$ (M$), applied to robust loss (RA-GO) or mean loss (M4)
Redesign	Cost $0.04S/n_{dec}$, extra delay $30/n_{dec}$ days, $n_{dec} = 10$ (amortized over reviews)
Stress world	Level 1.5: CVs and σ ×1.5, L_0 mean ×1.25, θ mean -0.1, readiness mean -0.15
Baseline default	Uniform intensity $x = 0.8$, no redesign (M1–M4)

D. F6: Mathematical Consistency, Binding Gates, and Feasibility

Property 1 (Consistent feasibility and halt safety): Assume $x_s^{gate} \in [0, 1]$. Then RA-GO v2 either returns HALT (when no gate-feasible active or status-quo option satisfies $\rho_\varepsilon[L] \leq \kappa_s$) or returns an option satisfying $m_a = \emptyset$, $\rho_\varepsilon[L] \leq \kappa_s$, and—whenever $a \neq a_{\mathrm{sq}}$—$x \geq x_s^{gate}$.

Property 2 (Robust monotonicity): $\rho_\varepsilon[f]$ is nondecreasing in ε, equals $\mathbb{E}_{\hat{P}}[f]$ at $\varepsilon = 0$, and never exceeds $\max_i f_i$. This follows from the nested ambiguity sets and Eq. (8).

Three independent checks support the consistency claim. First, the legal gates, the review-intensity floor, and the no-action option do not contradict one another: across 160 edge-case configurations built with tight caps and legacy statutory breaches, Property 1 held with 0 gate, x_s^{gate}, or cap violations, and all 88 configurations in which no admissible option existed returned HALT. The earlier v1 fallback violated the risk cap in 100% of tight-cap unviable tests ($0.12 \times \kappa_s$) and the intensity floor in 76.7% by reverting to a high-exposure status quo; the three-way partition prevents that fallback. Second, the gates

bind only because they are non-compensatory: unconstrained weighted scoring (M2) selects a gate-bypassing option in 45% of the base scenarios, whereas M1, M3, M4, and RA-GO never do. Third, the governance-frontier propositions of Section V-B that RA-GO operationalizes hold numerically: on 20,000 random parameter regimes (16,791 interior optima, 3,209 corner solutions with $x^* = 0$) we found 0 violations of Propositions 1–4, with a largest relative error of 7.5e-05 for the analytic derivative $dx^*/d\theta$ against a finite difference and 1.1e-03 for the envelope result. This is a consistency check of the algebra and the implementation, not an empirical result, and Proposition 5 can only be tested on process data (Section VIII-G).

On identical inputs the mean decision time per scenario is about 0.41 s for RA-GO against 0.2 ms for GAIM-DF, reflecting the dual solve in Eq. (8) for each option and metric. Min–max criteria normalization is set-dependent (removing Pareto-dominated options shifted the nominal-weight winner in 10 of 400 base runs), and $\mathbb{E}_{\mathrm{Dir}(cw_0)}[w^\top \tilde{z}] = w_0^\top \tilde{z}$ implies that Dirichlet sampling tracks nominal linear scoring closely; F5 is therefore a triage result, not a ranking result.

E. F4: KL-DRO Tail-Risk Protection Under Misspecification

Comparisons equalize the action space where noted, and OOD worlds are generated independently of RA-GO's training assumptions.

Equalized action space. In the base experiment, baselines M1–M4 were restricted to $g = 0$ and fixed $x = 0.8$ while RA-GO could choose $(g, x) \in \{0, 1\} \times [x_s^{gate}, 1]$. Giving GAIM-DF the same (g, x) candidate optimization and Pareto filter (M3+, Table XX, Panel A) lowers its stress-world expected loss from 1.57 to 1.44 M$ (55% of the gap to RA-GO at 1.32 M$). M3+ then records the lowest risk-neutral net-value regret (1.99 vs. 2.32 M$) and the lowest tail-adjusted regret at $n_{dec} = 10$ (1.46 vs. 1.64 M$). When redesign is charged to a single decision ($n_{dec} = 1$), restricted M3 beats RA-GO on both regret metrics. Under equalized actions, RA-GO's remaining advantage is concentrated in tail loss and risk-cap safety.

Structural OOD benchmark. We generated 30 non-template scenarios (6–8 Pareto-competing alternatives, random $n_{dec} \in \{1, 2, 5, 10\}$, 25% legacy statutory breaches) and evaluated decisions trained *solely* on RA-GO's nominal lognormal/exponential assumptions against four structural shifts (Table XX, Panel B): (OOD-1) multivariate Student-$t_{3.5}$ heavy tails with positive copula correlation ($\rho \in [0.50, 0.60]$) among cost, delay, and exposure; (OOD-2) thresholded review efficacy with reviewer fatigue above $x = 0.70$, violating concavity; (OOD-3) $M/M/1$ congestion delay $\propto u/(1-u)$ with stochastic redesign failure ($\Pr(\xi < 0) = 1 - R$); and (OOD-5) all shifts combined. Under OOD-5 the assumed redesign layer increases mean latency and risk-neutral regret relative to restricted GAIM-DF, while KL-DRO preserves tail safety: RA-GO v2 records 0.0% true risk-cap violations against 6.7% for M3+ and 26.7% for M1, and reduces Conditional Value-at-Risk $\mathrm{CVaR}_{0.95}(L)$ to 11.80 M$ against 14.11 for M3+ and 13.23 for restricted M3. Raising ε from 0 to 0.5 lowers stress-world p95 loss from 3.70 to 3.34 M$ (GAIM-DF: 4.11) with little change in regret.

Matched Shapley attribution. Averaging marginal effects over all $3! = 6$ layer orderings on strictly matched common non-halted scenario-seeds ($n = 289/300$ in OOD-5, which removes halt-composition selection bias; Table XXI) separates the layers cleanly. Friction redesign (F) drives $\mathrm{CVaR}_{0.95}$ reduction in the in-distribution stress world ($\phi = -0.51$ M$) but *increases* tail loss under OOD-5 ($\phi_{\mathrm{matched}} = +0.719$ M$; $+0.531$ M$ unmatched) once reorganization can fail and high intensity induces review fatigue, whereas KL-DRO (D) reduces tail risk within matched active cases ($\phi_{\mathrm{matched}} = -0.212$ M$ against -0.19 M$ in distribution) and across all cases including halts (-0.514 M$).

F. F5: Dirichlet EVPI Prioritizes Cases Needing Human Executive Review

Because $\mathbb{E}_{\mathrm{Dir}(cw_0)}[w^\top \tilde{z}] = w_0^\top \tilde{z}$, Dirichlet preference sampling is nearly neutral as a choice rule: its Shapley effect on loss is approximately zero and it changes the nominal-weight winner in only 1.4% of in-distribution stress runs and 6.4% of OOD-5 runs. Its measurable value lies elsewhere. Treating EVPI_s as a triage score and escalating the 25% of cases with the highest EVPI_s to human deliberation captures 1.56× the stakeholder-regret reduction per human review that random escalation achieves, and 1.08× when $1 - \max_o b_o^1$ is used instead (Fig. 14d). Under the dispersed stakeholder preferences modeled here, preference uncertainty is therefore informative about *which* decisions deserve scarce executive attention even though it is uninformative about which option to pick. The weights w_0 are assumed rather than elicited, and the regret reduction is computed inside the model.

G. F1–F3: Observational Approval-Latency Evidence and Falsification Tests

To test the observational implications of the governance-frontier model while accounting for confounding, clustering, and alternative case definitions, we analyzed TU Eindhoven's public payment and permit event logs (Fig. 15, Table XXII).

F1: approval latency. On 6,886 university non-travel payment and purchase cases (36,796 events; 0 unsorted timestamps) [154], 5,277 completed Admin-routed cases span 36 organizational units, and 35/36 units (99.9% of cases) contain both consolidated (Budget Owner = Supervisor, $n = 3,302$) and separate sequential ($n = 1,975$) sign-offs, so the comparison is made within rather than across units. Isolating the exact managerial stage (`Admin Approved` → `Final Supervisor Approved`) removes pre-submission drafting and post-approval batch triggers: mean stage delay is 2.71 d (median 1.09 d) under consolidated roles against 5.80 d (median 5.12 d) under separate roles (−53.4% unadjusted, cluster-bootstrap 95% CI $[43.4\%, 62.9\%]$, Cliff's $d = +0.538$). Controlling for organizational-unit, year-month, activity-code, and day-of-week fixed effects plus quadratic log-spend, a project indicator, and rework ($G = 36$ organizational clusters) yields $\hat{\beta} = -3.03$ d (95% CI $[-3.32, -2.74]$, -52.2%, $p_{\mathrm{cl}} = 1.4 \times 10^{-21}$; log-linear -52.1%, 95% CI $[-56.4\%, -47.4\%]$).

F2: overall process time. The association is not merely shifted to a later stage. Consolidation is associated with

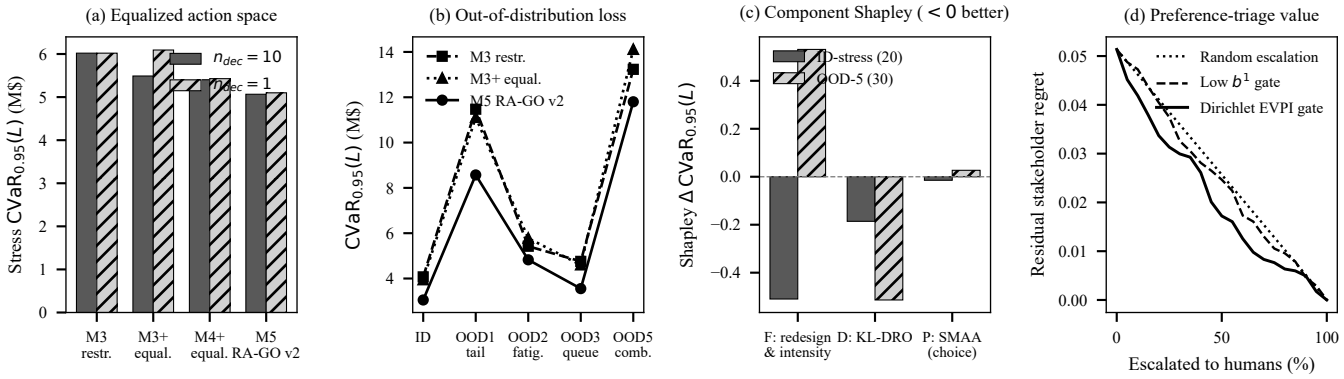


Fig. 14. Distribution-shift stress testing evidence for F4 and F5: (a) equalizing the (g, x) action space across baselines at $n_{dec} = 10$ vs. 1; (b) $\mathrm{CVaR}_{0.95}(L)$ across 30 non-template scenarios under structural and heavy-tail OOD shifts; (c) exact 2^3 Shapley contributions to $\mathrm{CVaR}_{0.95}(L)$, with friction redesign (F) reversing sign in OOD while KL-DRO (D) strengthens; (d) selective human-deliberation triage using Dirichlet EVPI_s vs. $1 - \max b^1$ and random escalation.

TABLE XX

EQUALIZED ACTION SPACE AND OUT-OF-DISTRIBUTION GENERALIZATION. PANEL A: CANDIDATE SPACE (a, g, x) AND PARETO FILTER EQUALIZED ACROSS GAIM-DF AND RA-GO V2 (20 BASE SCENARIOS, STRESS WORLD). PANEL B: DECISIONS TRAINED ON NOMINAL DISTRIBUTIONS EVALUATED ON 30 NON-TEMPLATE OOD SCENARIOS

Block / Regime	Method	$E[L]$	$\mathrm{CVaR}_{95}(L)$	$E[T]$ (d)	Regret(U)	Regret(U_{CVaR})	Pref. Regret	g=1/Halt %
Panel A: Action-space fairness audit on 20 base scenarios (10 seeds; last column = share choosing redesign g=1)								
n_{dec}=10 (amortized)	M3 GAIM-DF (restricted)	1.57	6.02	183	2.08	1.75	0.049	0
n_{dec}=10 (amortized)	M3+ GAIM-DF (equalized)	1.44	5.49	182	1.99	1.46	0.039	97
n_{dec}=10 (amortized)	M4+ GAIM-DF+cap (equalized)	1.41	5.40	174	2.27	1.71	0.052	100
n_{dec}=10 (amortized)	M5 RA-GO v2	1.32	5.07	178	2.32	1.64	0.055	100
n_{dec}=1 (one-off)	M3 GAIM-DF (restricted)	1.57	6.02	183	2.04	1.68	0.042	0
n_{dec}=1 (one-off)	M3+ GAIM-DF (equalized)	1.59	6.09	180	2.00	1.66	0.040	5
n_{dec}=1 (one-off)	M4+ GAIM-DF+cap (equalized)	1.42	5.43	191	2.92	2.34	0.060	62
n_{dec}=1 (one-off)	M5 RA-GO v2	1.33	5.10	193	2.88	2.18	0.063	54
Panel B: Structural & distributional OOD benchmark on 30 non-template scenarios (8 seeds; last column = true cap-violation % / halt %)								
In-distribution	M3 GAIM-DF (restricted)	1.48	4.07	157	2.43	2.12	0.063	0.0 / 0.0
In-distribution	M3+ GAIM-DF (equalized)	1.44	3.94	158	2.41	2.05	0.058	0.0 / 0.0
In-distribution	M4+ GAIM-DF+cap (equalized)	1.35	3.72	179	2.27	1.84	0.069	0.0 / 0.0
In-distribution	M5 RA-GO v2	1.11	3.05	182	2.18	1.65	0.070	0.0 / 3.3
OOD-1: $t_{3.5}$ tails + copula	M3 GAIM-DF (restricted)	1.75	11.47	180	1.43	1.69	0.063	0.0 / 0.0
OOD-1: $t_{3.5}$ tails + copula	M3+ GAIM-DF (equalized)	1.69	11.11	182	1.40	1.50	0.059	3.3 / 0.0
OOD-1: $t_{3.5}$ tails + copula	M4+ GAIM-DF+cap (equalized)	1.59	10.41	207	1.40	1.19	0.068	1.7 / 0.0
OOD-1: $t_{3.5}$ tails + copula	M5 RA-GO v2	1.31	8.57	211	1.35	0.73	0.069	0.0 / 3.3
OOD-2: Threshold + fatigue	M3 GAIM-DF (restricted)	1.70	5.42	157	2.25	1.77	0.059	0.0 / 0.0
OOD-2: Threshold + fatigue	M3+ GAIM-DF (equalized)	1.81	5.78	158	2.38	2.03	0.063	0.0 / 0.0
OOD-2: Threshold + fatigue	M4+ GAIM-DF+cap (equalized)	1.69	5.41	179	2.22	1.74	0.073	0.0 / 0.0
OOD-2: Threshold + fatigue	M5 RA-GO v2	1.51	4.83	183	2.21	1.66	0.078	0.0 / 3.3
OOD-3: $M/M/1$ + redesign fail.	M3 GAIM-DF (restricted)	1.48	4.76	172	2.52	2.06	0.063	0.0 / 0.0
OOD-3: $M/M/1$ + redesign fail.	M3+ GAIM-DF (equalized)	1.44	4.60	179	2.49	1.97	0.062	0.0 / 0.0
OOD-3: $M/M/1$ + redesign fail.	M4+ GAIM-DF+cap (equalized)	1.35	4.34	200	2.35	1.75	0.071	0.0 / 0.0
OOD-3: $M/M/1$ + redesign fail.	M5 RA-GO v2	1.11	3.55	206	2.30	1.57	0.073	0.0 / 3.3
OOD-5: All shocks combined	M3 GAIM-DF (restricted)	2.00	13.23	201	1.73	1.67	0.062	3.3 / 0.0
OOD-5: All shocks combined	M3+ GAIM-DF (equalized)	2.14	14.11	210	1.83	2.14	0.071	6.7 / 0.0
OOD-5: All shocks combined	M4+ GAIM-DF+cap (equalized)	2.00	13.15	235	1.77	1.67	0.078	5.8 / 0.0
OOD-5: All shocks combined	M5 RA-GO v2	1.78	11.80	243	1.87	1.65	0.083	0.0 / 3.3

"Restricted" fixes g=0 and x=0.8; "equalized" grants the same (g, x) candidate optimization and Pareto filter as RA-GO v2. OOD-1 uses Student-$t_{3.5}$ heavy tails with copula correlation $\rho \in [0.50, 0.60]$; OOD-2 thresholded review efficacy with reviewer fatigue above $x = 0.70$; OOD-3 $M/M/1$ congestion with stochastic redesign failure; OOD-5 all shocks combined. $U = V - C - L - \lambda T$ (risk-neutral net value, M$);
$U_{\mathrm{CVaR}} = V - C - \lambda T - 0.5E[L] - 0.5\,\mathrm{CVaR}_{0.95}(L)$ (tail-risk-sensitive net value); Pref. Regret: normalized stakeholder utility loss under dispersed true weights $w_{\mathrm{true}} \sim \mathrm{Dir}(5w_0)$. In OOD-4 tight-cap safety tests ($0.12 \times \kappa_s$), the original v1 fallback violated the risk cap in 100% of cases and $x \geq x_s^{gate}$ in 76.7% by reverting to high-exposure status quo; corrected RA-GO v2 triggered Mandatory Halt in 100% with 0% violations.

-41.9% FE-adjusted delay on `First Submit` → `Final Approval` ($\hat{\beta} = -3.47$ d, 95% CI $[-4.35, -2.59]$) and, on the full `Submit` → `Request Payment` window that includes the post-approval batch trigger, with -34.9% unadjusted (7.87 vs. 12.09 d, 95% CI $[24.1\%, 43.5\%]$), -31.7% after the same fixed effects and controls ($\hat{\beta} = -3.83$ d, 95% CI $[-4.93, -2.74]$), and -35.3% under exact 335-stratum matching. The adjusted estimate holds in every spend quartile (Fig. 15b) and on clean first-pass approvals (-38.5%, $n = 4,777$), so it is not produced by a few large or reworked cases. On the Permit Log ($n = 7,002$ complete cases, 28 organizational units) [154], the converse arrangement shows the same direction: adding a fourth-tier `DIRECTOR` approval adds $+4.38$ d after organizational, month, and log-budget

TABLE XXI
EXACT 2^3 FACTORIAL SHAPLEY DECOMPOSITION OF RA-GO'S THREE LAYERS (AVERAGED OVER ALL $3! = 6$ ORDERINGS; NEGATIVE ΔL, $\Delta\mathrm{CVaR}_{95}$, AND ΔREGRET IMPROVE PERFORMANCE). OOD-5 COLUMNS REPORT STRICTLY MATCHED COMMON NON-HALTED SCENARIO-SEEDS ($n = 289/300$) TO ELIMINATE HALT-COMPOSITION SELECTION BIAS, WITH UNMATCHED NON-HALTED VALUES IN PARENTHESES

	Base 20 Scenarios (ID-Stress, $n = 200/200$)				30 OOD-5 Scenarios (Matched $n = 289/300$; Unmatched in Paren.)			
Component	$\phi(E[L])$	$\phi(\mathrm{CVaR}_{95})$	$\phi(E[T])$	$\phi(\mathrm{Reg}_{\mathrm{CVaR}})$	$\phi(E[L])$	$\phi(\mathrm{CVaR}_{95})$	$\phi(E[T])$	$\phi(\mathrm{Reg}_{\mathrm{CVaR}})$
F: Friction redesign + intensity	-0.134	-0.510	+5.2	-0.153	+0.112	+0.719 (+0.53)	+14.3	+0.527
D: KL-ball DRO ($\varepsilon = 0.19$)	-0.048	-0.186	+3.2	-0.081	-0.032	-0.212 (-0.51)	+11.6	-0.013
P: Dirichlet SMAA (choice only)	-0.004	-0.015	+1.3	+0.015	+0.002	+0.027 (+0.03)	-1.3	+0.022

Shapley efficiency holds exactly ($\sum_i \phi_i = v(1,1,1) - v(0,0,0)$ to $< 10^{-4}$). Conditioning on identical non-halted pairs ($n = 289/300$) isolates within-case DRO tail reduction ($\phi = -0.212$ M$) from the additional tail-avoidance of halting the 11 riskiest scenario-seeds (-0.514 M$ unmatched). Pure Dirichlet SMAA alters only 1.4% (ID) and 6.4% (OOD-5) of nominal choices ($\mathbb{E}[w] = w_0$); its operational value is selective EVPI_s human triage (1.56× random escalation at a 25% review budget).

controls (95% CI $[+2.56, +6.19]$, $p_{\mathrm{cl}} = 3.5 \times 10^{-5}$, Cliff's $d = +0.402$).

F3: falsification tests. Role consolidation has no mechanism to act on stages it does not touch, so each such stage is a falsification opportunity. In all three the estimate is statistically indistinguishable from zero (Fig. 15a): upstream `Submit` → `Admin` ($\hat{\beta} = -0.48$ d, $p_{\mathrm{cl}} = 0.24$), the post-approval batch trigger ($\hat{\beta} = -0.52$ d, $p_{\mathrm{cl}} = 0.21$), and bank clearing ($\hat{\beta} = -0.06$ d, $p_{\mathrm{cl}} = 0.41$). A unit-level or seasonal confounder that produced F1 and F2 artifactually would be expected to move these windows as well; the null estimates therefore locate the association in the managerial stage. Role overlap remains observational rather than randomized.

The accompanying trade-off. Across all $5,349$ Admin-approved cases, separate sequential Budget-Owner-plus-Supervisor review rejects 5.61% $(113/2,016)$ of requests after Admin approval against 1.92% $(64/3,333)$ under consolidated review (Fisher $p = 1.2 \times 10^{-12}$; Fig. 15c), and overall rejection rises to 22.5% in top-quartile spend against 12.3–14.0% in Q1–Q3. This is a 3.69-percentage-point reduction in the observed post-Admin rejection rate (5.61% to 1.92%), consistent with exposure-tiered rather than blanket consolidation (Proposition 4). The logs cover general university expenses and travel permits, not security or privacy reviews, and carry no dollar-loss labels.

H. Scope of the Evidence

Table XXIII records the evidence class and claim boundary for each finding; Table XXIV summarizes the corresponding validity threats.

EVPI_s to triage contested options to human committees (F5). Non-compensatory screening in MCDA, Kullback–Leibler DRO and its log-exponential dual [173]–[176], Dirichlet rank-acceptability indices identical to SMAA-2 (Stochastic Multicriteria Acceptability Analysis) [177], [178], Pareto filtering [182], and robust supplier-selection formulations [183], [184] are established methods reused within that formulation.

IX. LIMITATIONS

Fig. 16 shows three modernization pathways (regulation-encoded process redesign, governance restructuring, and AI-assisted decision support) across assess, pilot, scale, and sustain phases with four human checkpoints.

X. EXECUTIVE GUIDELINES

Table XXV presents twelve guidelines for executives governing modernization decisions, each tied to evidence strength and to retained human authority.

A. Illustrative Executive Decision Path

Consider an organization-independent case in which an executive must choose among retaining a legacy system, modernizing it incrementally, or replacing it with an AI-enabled alternative. Under GAIM-DF, intake records scope and decision rights; Gate H1 confirms applicable privacy, security, accessibility, and procurement constraints for each option; infeasible options leave the set F; remaining options are compared on expected benefits, cost and resource load, residual risk, and evidence confidence, with uncertainty intervals reported when scores overlap; Gate H2 records risk acceptance by the accountable owner; and Gate H3 is the authorized human decision. RA-GO v2, when used, applies non-compensatory gates, risk-adaptive evaluation over feasible options, and EVPI_s triage for scarce human review; if no admissible option remains, the output is HALT.

Supplementary portfolio datasets are descriptive only: in the University of Washington portfolio [155] oversight tier scales with budget ($\rho = 0.638$, 95% CI $[0.117, 0.914]$, Holm-adjusted permutation $p = 0.041$, $n = 11$), but the $18/25$ threshold is UW's own Red-status definition, so an exact rubric match is retrospective consistency rather than prediction, and Leave-One-Out cross-validation gives 75.0% balanced accuracy ($\mathrm{MCC} = 0.671$) on 11 projects with two failures (Fig. 15d); in Oregon [179] the active Stage-Gate number tracks lifecycle phase rather than budget ($\rho = 0.033$, $p = 0.74$); and in Mesa [180] 36.0% and 61.0% missing flags widen complete-case adherence to Manski bounds of $[58.0\%, 94.0\%]$

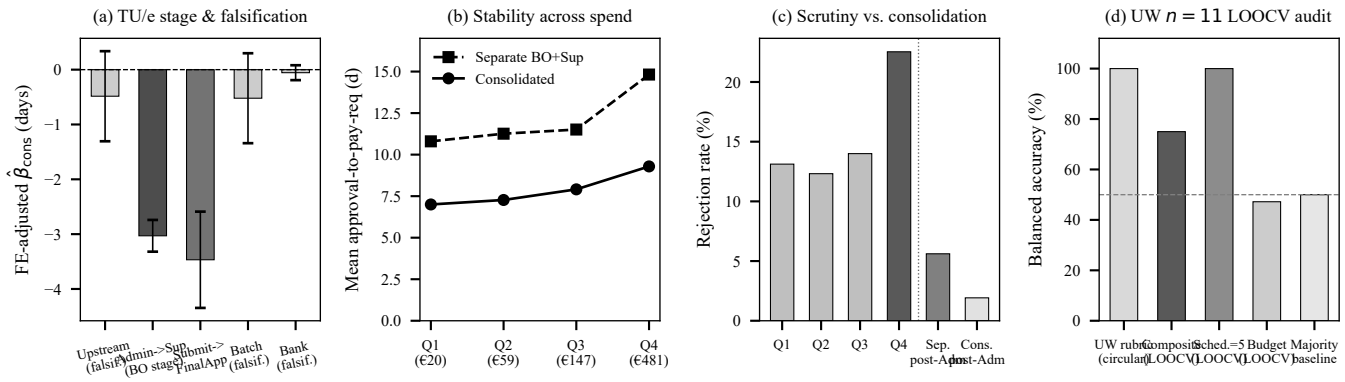


Fig. 15. Observational evidence for F1–F3: (a) TU Eindhoven FE-adjusted consolidation coefficients ($\hat{\beta}_{\text{cons}}$, 95% cluster CIs) across approval stages against the three null falsification stages; (b) mean approval delay across spend quartiles; (c) spend-quartile rejection rates and the post-Admin speed–scrutiny trade-off (5.61% separate vs. 1.92% consolidated); (d) University of Washington ($n = 11$) rubric circularity against Leave-One-Out cross-validated balanced accuracy.

TABLE XXII
OBSERVATIONAL EVIDENCE FOR F1–F3 ON PUBLIC UNIVERSITY EVENT LOGS

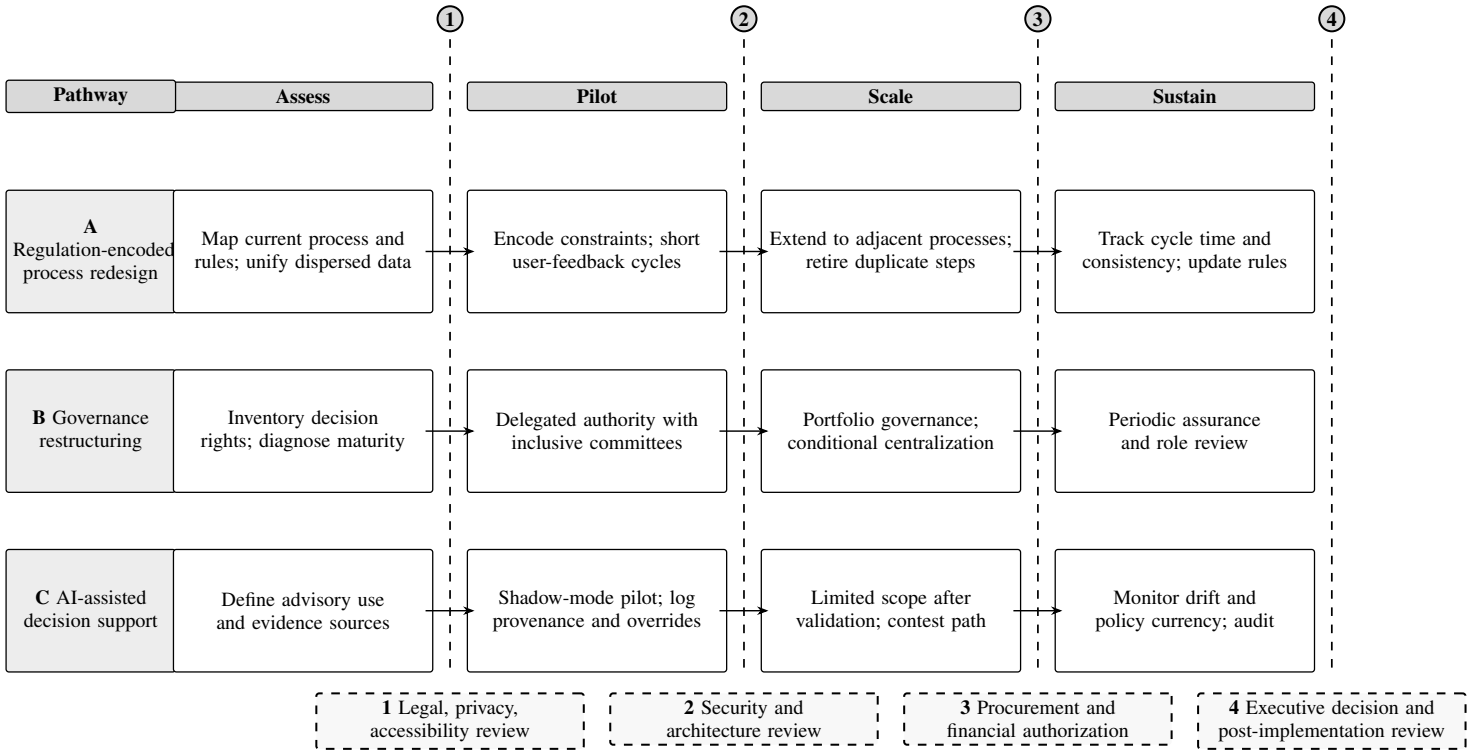
Dataset & Provenance	Sample & Units	Confounder-Adjusted Findings, Effect Sizes, & Falsification Checks	Trade-Offs & Inferential Boundaries
TU Eindhoven Payment Approval Log [154] (CC BY-NC 4.0; MD5 verified; 0 unsorted timestamps)	$6,886$ cases ($36,796$ events); $5,277$ completed Admin-routed cases across 36 org units ($35/36$ units contain both regimes, 99.9% overlap)	**F1 (stage-isolated latency):** In the post-Admin managerial stage (`Admin Approved` → `Final Supervisor Approved`), consolidated Budget Owner = Supervisor ($n = 3,302$) averages **2.71 d** (median 1.09 d) vs. **5.80 d** (median 5.12 d) under separate sequential sign-offs ($n = 1,975$; unadjusted -53.4%, cluster-bootstrap 95% CI $[43.4\%, 62.9\%]$, Cliff's $d = +0.538$). With Org-Unit, Year-Month, Activity, and Day-of-Week fixed effects plus quadratic log-spend, project, and rework controls ($G = 36$ clusters), adjusted $\hat{\beta} = \mathbf{-3.03}$ d (95% CI $[-3.32, -2.74]$, $\mathbf{-52.2}\%$, $p_{\text{cl}} = 1.4 \times 10^{-21}$; log-linear -52.1%, 95% CI $[-56.4\%, -47.4\%]$). **F2 (overall process time):** `Submit` → `Request Payment` (incl. post-approval batch trigger) is -34.9% unadjusted (7.87 vs. 12.09 d, 95% CI $[24.1\%, 43.5\%]$), $-\mathbf{31.7}\%$ FE-adjusted ($\hat{\beta} = -3.83$ d, 95% CI $[-4.93, -2.74]$), and -35.3% under exact 335-stratum matching; `Submit` → `Final Approval` is -41.9% ($\hat{\beta} = -3.47$ d, 95% CI $[-4.35, -2.59]$); stable in all four spend quartiles and on clean first-pass cases (-38.5%, $n = 4,777$). **F3 (falsification, null as predicted):** upstream `Submit` → `Admin` ($\hat{\beta} = -0.48$ d, $p_{\text{cl}} = 0.24$), batch trigger ($\hat{\beta} = -0.52$ d, $p_{\text{cl}} = 0.21$), bank clearing ($\hat{\beta} = -0.06$ d, $p_{\text{cl}} = 0.41$).	**Speed–scrutiny trade-off:** across all $5,349$ Admin-approved cases, separate sequential review rejects **5.61%** ($113/2,016$) after Admin approval vs. **1.92%** ($64/3,333$) under consolidated review (Fisher $p = 1.2 \times 10^{-12}$), and Q4 spend shows 22.5% overall rejection vs. 12.3–14.0% in Q1–Q3—so consolidation should be exposure-tiered rather than blanket. Role overlap is observational, not randomized; the log covers general university expenses and hardware, not FERPA or cybersecurity reviews.
TU Eindhoven Travel Permit Log [154] (CC BY-NC 4.0; MD5 verified)	$7,065$ cases ($86,581$ events); $7,002$ with complete approval timestamps (28 org units)	Converse test of the same mechanism: adding a 4th executive sign-off (`DIRECTOR`, $n = 676$ vs. $6,326$) raises mean approval latency from 3.64 d (median 1.68 d) to 8.23 d (median 4.96 d; Cliff's $d = +0.402$). With Org-Unit FE, Year-Month FE, and log budget ($G = 28$ clusters), adjusted $\hat{\beta} = +\mathbf{4.38}$ d (95% CI $[+2.56, +6.19]$, $p_{\text{cl}} = 3.5 \times 10^{-5}$).	Quantifies marginal executive-tier handoff latency net of trip budget and departmental fixed effects; travel permits are lower-exposure decisions than enterprise modernization or AI-adoption reviews.
Portfolio boundary datasets: UW OCIO [155], Oregon EIS [179], Mesa [180], Texas QAT [153]	$n = 11/14$ UW projects; 115 Oregon projects; 400 Mesa projects; 59 Texas major projects	Reported only as boundary evidence on exposure-tiered oversight, including null and wide-bound results: UW oversight tier scales with budget ($\rho = 0.638$, 95% CI $[0.117, 0.914]$, Holm permutation $p = 0.041$), but $18/25$ is UW's own Red-status definition, so rubric matching is retrospective consistency and LOOCV balanced accuracy is 75.0% (MCC $= 0.671$); Oregon Stage-Gate number tracks lifecycle phase, not budget ($\rho = 0.033$, $p = 0.74$, null); Mesa missingness (36.0%, 61.0%) widens adherence to Manski bounds $[58.0\%, 94.0\%]$ and $[38.2\%, 99.2\%]$; Texas QAT reports $34/59$ (57.6%, 95% CI $[44.1\%, 70.4\%]$) projects within 10% of initial budget and schedule.	These datasets support none of F1–F6 on their own: samples are small or incomplete, ratings are contemporaneous rather than ex-ante, and no predictive-accuracy claim is made from them.

TABLE XXIII
SOURCE-TO-CLAIM EVIDENCE MATRIX FOR THE SIX PRINCIPAL FINDINGS: EVIDENCE CLASS, EFFECT SIZE WITH UNCERTAINTY, TRADE-OFF OR BOUNDARY, AND INTERPRETATION

Finding	Evidence Class	Primary Verifiable Evidence, Effect Size, & Uncertainty	Trade-Off / Boundary	Interpretation
F1. -52.2% approval latency: coordination of sign-off roles	**Observationally Supported** (adjusted association, not causal)	TU/e payment log [154], $n = 5,277$ completed cases, 36 org units, $35/36$ units containing both regimes. Stage-isolated `Admin→Final Supervisor`: 2.71 vs. 5.80 d (-53.4% unadjusted, cluster-bootstrap 95% CI $[43.4\%, 62.9\%]$, Cliff's $d = +0.538$); with org, year-month, activity, day-of-week FE plus quadratic log-spend, project, and rework controls ($G = 36$ clusters), $\hat{\beta} = -3.03$ d (95% CI $[-3.32, -2.74]$, -52.2%, $p_{\text{cl}} = 1.4 \times 10^{-21}$; log-linear -52.1%, 95% CI $[-56.4\%, -47.4\%]$).	Not randomized; logs carry no dollar-loss label, so the coordination gain is measured in days, not in risk avoided.	Supported as a confounder-adjusted association; not a causal estimate
F2. -31.7% overall process time after adjustment	**Observationally Supported** (adjusted association, not causal)	Same log, full `Submit→Request Payment` window: -34.9% unadjusted (7.87 vs. 12.09 d, 95% CI $[24.1\%, 43.5\%]$), -31.7% FE-adjusted ($\hat{\beta} = -3.83$ d, 95% CI $[-4.93, -2.74]$), -35.3% under exact 335-stratum matching, -41.9% on `Submit→Final Approval` ($\hat{\beta} = -3.47$ d, 95% CI $[-4.35, -2.59]$); holds in all four spend quartiles and on clean first-pass cases (-38.5%, $n = 4,777$). Permit log: a fourth `DIRECTOR` tier adds $+4.38$ d (95% CI $[+2.56, +6.19]$).	The unadjusted -34.9% window includes a post-approval batch trigger; the adjusted -31.7% is the synthesis figure.	Supported as an adjusted association across case definitions
F3. Falsification tests support the latency findings	**Statistical robustness check** on observational data	Three stages that role consolidation cannot act on are null as predicted: upstream `Submit→Admin` ($\hat{\beta} = -0.48$ d, $p_{\text{cl}} = 0.24$), post-approval batch trigger ($\hat{\beta} = -0.52$ d, $p_{\text{cl}} = 0.21$), bank clearing ($\hat{\beta} = -0.06$ d, $p_{\text{cl}} = 0.41$). Negative controls use the same specification, clusters, and sample.	Nullity makes a purely spurious reading less plausible but cannot exclude a confounder acting only on the managerial stage.	Supported as a credibility check; causality remains unproven
F4. KL-DRO gives tail-risk protection under misspecification	**Scenario-Based Stress Testing** (constructed ID and structural OOD)	Across 30 non-template OOD scenarios, RA-GO v2 records 0.0% true risk-cap violations vs. 6.7% (equalized M3+) and 26.7% (M1), cutting $\text{CVaR}_{0.95}(L)$ to 11.80 M\$ (M3+: 14.11). Matched exact 2^3 Shapley ($n = 289/300$) attributes $\phi_{\text{matched}} = -0.212$ M\$ to KL-DRO (ID: -0.19; -0.514 M\$ including halts) while friction redesign reverses to $+0.719$ M\$.	Equalizing (g, x) explains 55% of the base loss gap and M3+ wins on risk-neutral regret; magnitudes depend on the constructed distributions and $\varepsilon = 0.19$.	Supported under scenario-based stress testing for tail risk and cap safety
F5. Dirichlet EVPI_s prioritizes cases for human executive review	**Mathematically Analyzed + Numerical Experiments**	Because $\mathbb{E}_{\text{Dir}(cw_0)}[w^\top \tilde{z}] = w_0^\top \tilde{z}$, pure SMAA alters only 1.4% (ID) and 6.4% (OOD-5) of nominal choices ($\phi \approx 0$). Escalating the top 25% of EVPI_s cases yields $1.56\times$ the stakeholder-regret reduction per review of random escalation ($1.08\times$ using $1 - \max_o b_o^1$).	Weights w_0 are assumed, and regret reduction is computed inside the model; committee behavior is not observed.	Supported in numerical experiments as triage; not tested with elicited human weights
F6. Mathematical consistency and mandatory governance gates	**Mathematically Proven** (by construction) + **Numerically Verified**	Property 1 enforces $m_a = \emptyset$, $\rho_\varepsilon[L] \le \kappa_s$, and $x \ge x_s^{gate}$: $0/160$ edge-case violations and $88/88$ unviable cases halted, against 100% cap violations for the v1 status-quo fallback. Unconstrained scoring (M2) bypasses gates in 45% of scenarios. Propositions 1–4 hold on 20,000 random regimes with 0 violations (derivative error $\le 7.5e-05$). Mean decision time ≈ 0.41 s per scenario.	An architectural and analytic invariant, not an empirical discovery; no IIA property is claimed, and no institutional policy was reviewed.	Supported as a formal invariant and feasibility result; human behavioral effects not observed

TABLE XXIV
THREATS TO VALIDITY BOUNDING THE SIX FINDINGS

Threat	Status
Observational design (F1–F3)	Role overlap is not randomized; estimates are confounder-adjusted associations. The observed post-Admin rejection rate falls by 3.69 percentage points, from 5.61% under separate sequential review to 1.92% under consolidated review. Logs cover payments and travel permits, not security, privacy, or AI reviews, and carry no dollar-loss labels.
Constructed parameters (F4–F5)	Modeled costs, exposures, delays, readiness values, and weights are assumed; magnitudes are conditional on ε, n_{eff}, and c.
Action-space & n_{dec} bias	55% of the base loss gap came from restricting M3 to fixed $(g{=}0, x{=}0.8)$; redesign fails under $n_{dec} = 1$ or OOD reorganization shocks.
SMAA & normalization	Pure Dirichlet SMAA alters $<7\%$ of nominal choices ($\mathbb{E}[w] = w_0$) and is used for EVPI_s triage rather than standalone ranking. Min–max normalization is set-dependent and shifted the nominal-weight winner in 2.5% of base runs.
Legal scope (F6)	Gates are binding inside the modeled constraint set and are drafted from verified primary sources; institutional policy review and legal compliance are not established here.



Numbered circles and dashed lines mark human oversight checkpoints between phases; each pathway passes the same four checkpoints. Pathways may run in parallel; C depends on A or B having made rules and data explicit.

Fig. 16. Illustrative modernization pathways by phase with four human oversight checkpoints.

TABLE XXV
EXECUTIVE GUIDELINES DERIVED FROM THE EVIDENCE

ID	Guideline	Evidence basis (strength)	Authority that remains human
E1	Document decision rights and escalation paths for each approval step before changing any step.	Fragmented ownership and unit silos [93], [157] (low).	Governance chair; accountable executive.
E2	Use conditional rather than blanket centralization: centralize where infrastructure is heterogeneous or risk is high, and delegate execution within guardrails.	[3], [81], [84] (moderate for security; low otherwise).	Accountable executive; CIO with CISO.
E3	Treat required safeguards as hard constraints and measure the time spent in each review stage to find avoidable delay.	Hierarchy supports legal certainty [4]; lengthy approval stages [102] (low).	Legal counsel; privacy officer; CISO.
E4	Reform procurement paths to be compliant and early in the process, and write service-level and security terms into contracts.	[87], [103], [141], [181] (low).	Procurement officer; general counsel; CFO.
E5	Involve affected staff in design and communicate the reason and pace of change, since resistance is reported as a response to design.	[6], [7], [93] (low).	Accountable executive; workforce representatives.
E6	Assess readiness and governance maturity first, and treat low maturity as a prerequisite issue, not as resistance.	[98], [124] (low).	CIO; accountable executive.
E7	Govern modernization at portfolio level, not project by project.	[127], [128] (low).	Accountable executive; CIO.
E8	Use AI only as advisory support with evidence provenance, effective dates, logged overrides, and a contest path; do not delegate approval or rejection.	[19], [105], [109], [132] (very low).	Authorized approvers named per constraint.
E9	Do not treat an AI vendor or model selection as establishing legal compliance; obtain legal and security review of each use.	Accountability and assurance requirements [34], [61], [132] (analytical).	Legal counsel; CISO; privacy officer.
E10	Pilot any new decision-support process in shadow mode and report results with uncertainty before relying on it.	Absence of independent evaluation (Table X, G4).	Executive sponsor; authorized reviewers.
E11	Require, for any major legacy-system retirement, a documented plan with milestones, a description of the work, and the planned disposition of the legacy system.	Leading-practice elements reported by GAO for Federal systems [1] (transferability untested).	CIO; executive sponsor.
E12	Maintain a catalogue of hard constraints (privacy, accessibility, procurement, security) with sources and effective dates, and have counsel confirm applicability before each approval.	Table XII; consulted statutory and regulatory texts (applicability is a legal determination).	General counsel; privacy officer; CISO.

EVPI_s to triage contested options to human committees (F5). Non-compensatory screening in MCDA, Kullback–Leibler DRO and its log-exponential dual [173]–[176], Dirichlet rank-acceptability indices identical to SMAA-2 (Stochastic Multicriteria Acceptability Analysis) [177], [178], Pareto filtering [182], and robust supplier-selection formulations [183], [184] are established methods reused within that formulation.

X. LIMITATIONS

Threats to validity for F1–F6 are recorded in Tables XXIII and XXIV and are not restated here. Manuscript-level limits outside those tables: native Scopus, Web of Science, IEEE Xplore, and ACM portals were not queried; 53 of 129 reports sought for retrieval were not obtained as verified English full texts; screening and appraisal used one reviewer; and the review was not registered. Evaluations A, B, and D are deterministic rule-based implementations; Evaluation C uses one 8B language model, one seed, and 138 cases with rule-derived D-labels. The Hybrid in Section VII-G is proposed without measurements.

XI. CONCLUSION

The central contributions are GAIM-DF, a human-gated advisory process for executive modernization decisions, and RA-GO v2, which adds mandatory gates, risk-adaptive evaluation, and human-review triage. An optional multi-expert coordination extension, evaluated in four comparisons, characterizes conditions, trade-offs, and failure modes relative to a strong Single-Agent baseline and to RA-GO alone. A Hybrid architecture is proposed for future evaluation.

Across 13,287 uncapped records (11,265 unique; 129 sought for retrieval; 53 not retrieved; 69 full texts), barriers cluster around governance capacity (47), legacy and technical fragmentation (29), structural approval chains (25), and cultural and human factors (21). Privacy, security, accessibility, and procurement mandates remain non-compensatory human gates. Evidence classes and principal quantitative findings for F1–F6 are recorded in Table XXIII.

ACKNOWLEDGMENT

The authors thank the University of Houston's University Information Technology (UIT) division for its support.

FUNDING

This research received no external funding.

CONFLICTS OF INTEREST

The authors declare no conflicts of interest.

DATA AVAILABILITY

The publicly available BPI Challenge 2019 and 2020 datasets were used alongside computationally generated scenarios for controlled evaluation. Dataset sources are cited in the manuscript. Experimental artifacts are available from the corresponding author upon reasonable request. The search log, screening records, extraction tables, and scripts are inventoried in the project archive described in Appendix B.

REFERENCES

[1] U.S. Government Accountability Office, "Information technology: Agencies need to plan for modernizing critical decades-old legacy systems," U.S. GAO, Tech. Rep. GAO-25-107795, 2025. [Online]. Available: https://www.gao.gov/products/gao-25-107795

[2] ——, "High-risk series: Critical actions needed to urgently address IT acquisition and management challenges," U.S. GAO, Tech. Rep. GAO-25-107852, 2025. [Online]. Available: https://www.gao.gov/products/gao-25-107852

[3] S. S. Furusa and M. Phahlane, "From centralisation to regulated contestability: A multivocal policy synthesis of adaptive ict procurement governance for the state information technology agency in south africa," *Computers*, vol. 15, no. 11, p. 776, 2026.

[4] S. C. Sakinah and L. L. Fuadah, "The evolution of public sector management: A systematic literature review of traditional bureaucracy to modern governance," International Journal of Education Vocational and Social Science, 2026.

[5] B. Dunican and C. A. Gearin, "Exploring resistance to change and intolerance to ambiguity in higher education institutions," International Journal of New Trends in Arts, Sports & Science Education (IJTASE), 2018.

[6] E. Karadag and R. Yilmaz, "Confronting resistance to change: German higher education in times of rapid transformation," *Journal of Pedagogical Sociology and Psychology*, 2026.

[7] Y. F. Aldhafiri, "Coerced transformation: Vision 2030 as an institutional pressure field and the organisational dynamics of imposed digital change in saudi public sector," *Journal of Intelligent Decision Making and Information Science*, vol. 3, no. 6s, pp. 743–759, 2026.

[8] A. Rah, "Why digital transformation fails: Human resistance, legacy systems, and IT modernization challenges," ResearchGate preprint, 2026. [Online]. Available: https://doi.org/10.13140/RG.2.2.36366.11848

[9] ——, "Governance, human resistance, and cybersecurity challenges in digital transformation," ResearchGate preprint, 2026. [Online]. Available: https://doi.org/10.13140/RG.2.2.19233.36964

[10] M. Khouja, I. B. Rodriguez, Y. Ben Halima, and S. Moalla, "It governance in higher education institutions," *International Journal of Human Capital and Information Technology Professionals*, vol. 9, no. 2, pp. 52–67, 2018.

[11] M. Thomran, K. Alshammari, M. Grada, R. Alobaid, H. M.A. Khojally, A. S. Alshebami, H. Khandakar, and H. M.S. Ahmed, "Mapping strategies: investigating it governance approaches in higher education," *Global Knowledge, Memory and Communication*, 2025.

[12] Y. Magdalena, F. L. Gaol, N. Legowo, and B. Soewito, "The role and impact of it governance in higher education: A systematic literature review," in *2026 6th International Conference on Innovative Research in Applied Science, Engineering and Technology (IRASET)*, 2026, pp. 1–5.

[13] B. R. Aditya, R. Ferdiana, and S. S. Kusumawardani, "Categories for barriers to digital transformation in higher education: An analysis based on literature," *International Journal of Information and Education Technology*, vol. 11, no. 12, pp. 658–664, 2021.

[14] A. Rah, "Digital transformation beyond technology: Governance, human resistance, and cybersecurity," Preprints.org, 2026. [Online]. Available: https://doi.org/10.20944/preprints202605.1423.v1

[15] D. W. Putri, B. Riyanto, D. Muris, S. Xinyuan, C. Hantang, Y. Yucui, and A. S. Nagy, "Ai adoption in asean higher education as a digital governance challenge: A scoping review," *JeDEM - eJournal of eDemocracy and Open Government*, vol. 18, no. 2, pp. 192–223, 2026.

[16] A. Sigfrids, M. Nieminen, J. Leikas, and P. Pikkuaho, "How should public administrations foster the ethical development and use of artificial intelligence? a review of proposals for developing governance of ai," *Frontiers in Human Dynamics*, vol. 4, 2022.

[17] M. G. S. Gomes, A. S. M. Ramos, and K. d. Oliveira, "Guidelines and frameworks of governance for the adoption of generative artificial intelligence in higher education institutions: a scoping review," *Revista de Administração da UFSM*, vol. 18, no. 4, p. e1, 2025.

[18] A. Rah and D. W. Johnson, "Enterprise AI at the decision point: A U.S.-focused systematic review of productivity, ROI, workforce change, and governance," SSRN preprint, 2026. [Online]. Available: https://papers.ssrn.com/abstract=7517919

[19] ——, "Beyond AI productivity: Enterprise AI systems, human–AI responsibility, and the future of work," Research Square preprint, 2026. [Online]. Available: https://doi.org/10.21203/rs.3.rs-11670337/v1

[20] P. J. DiMaggio and W. W. Powell, "The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields," *American Sociological Review*, vol. 48, no. 2, p. 147, 1983.

[21] J. Mahoney, "Path dependence in historical sociology," *Theory and Society*, vol. 29, no. 4, pp. 507–548, 2000.

[22] M. T. Hannan and J. Freeman, "Structural inertia and organizational change," *American Sociological Review*, vol. 49, no. 2, p. 149, 1984.

[23] K. E. Weick, "Educational organizations as loosely coupled systems," *Administrative Science Quarterly*, vol. 21, no. 1, p. 1, 1976.

[24] M. D. Cohen, J. G. March, and J. P. Olsen, "A garbage can model of organizational choice," *Administrative Science Quarterly*, vol. 17, no. 1, p. 1, 1972.

[25] F. D. Davis, "Perceived usefulness, perceived ease of use, and user acceptance of information technology," *MIS Quarterly*, vol. 13, no. 3, pp. 319–340, 1989.

[26] V. Venkatesh, M. G. Morris, G. B. Davis, and F. D. Davis, "User acceptance of information technology: Toward a unified view1," *MIS Quarterly*, vol. 27, no. 3, pp. 425–478, 2003.

[27] D. J. Teece, G. Pisano, and A. Shuen, "Dynamic capabilities and strategic management," *Strategic Management Journal*, vol. 18, no. 7, pp. 509–533, 1997.

[28] G. Vial, "Understanding digital transformation: A review and a research agenda," *The Journal of Strategic Information Systems*, vol. 28, no. 2, pp. 118–144, 2019.

[29] A. Rah, "A governance-oriented analytical research paper on digital transformation and cybersecurity," Research Square preprint, 2026. [Online]. Available: https://doi.org/10.21203/rs.3.rs-9734788/v1

[30] J. M. Bryson, B. C. Crosby, and L. Bloomberg, "Public value governance: Moving beyond traditional public administration and the new public management," *Public Administration Review*, vol. 74, no. 4, pp. 445–456, 2014.

[31] E. L. Trist and K. W. Bamforth, "Some social and psychological consequences of the longwall method of coal-getting," *Human Relations*, vol. 4, no. 1, pp. 3–38, 1951.

[32] ISACA, *COBIT 2019 Framework: Governance and Management Objectives*. ISACA, 2018.

[33] National Institute of Standards and Technology, "The NIST cybersecurity framework (CSF) 2.0," National Institute of Standards and Technology, Tech. Rep. NIST CSWP 29, 2024.

[34] E. Tabassi, "Artificial intelligence risk management framework (AI RMF 1.0)," National Institute of Standards and Technology, Tech. Rep. NIST AI 100-1, 2023.

[35] National Institute of Standards and Technology, "Artificial intelligence risk management framework: Generative artificial intelligence profile," National Institute of Standards and Technology, Tech. Rep. NIST AI 600-1, 2024.

[36] Joint Task Force, "Security and privacy controls for information systems and organizations," National Institute of Standards and Technology, Tech. Rep. NIST SP 800-53 Rev. 5, 2020.

[37] J. Boyens, A. Smith, N. Bartol, K. Winkler, A. Holbrook, and M. Fallon, "Cybersecurity supply chain risk management practices for systems and organizations," National Institute of Standards and Technology, Tech. Rep. NIST SP 800-161 Rev. 1, 2022.

[38] U.S. Government Accountability Office, "Artificial intelligence: An accountability framework for federal agencies and other entities," U.S. GAO, Tech. Rep. GAO-21-519SP, 2021. [Online]. Available: https://www.gao.gov/products/gao-21-519sp

[39] A. Rah, "Cybersecurity risk controls in IT organizations," ResearchGate preprint, 2026. [Online]. Available: https://doi.org/10.13140/RG.2.2.32591.24488

[40] ——, "Auditing operational data records to improve trust and risk assessments," ResearchGate preprint, 2026. [Online]. Available: https://doi.org/10.13140/RG.2.2.25041.49768

[41] ——, "Reducing insider risk in identity recovery through zero-exposure governance," ResearchGate preprint, 2026. [Online]. Available: https://doi.org/10.13140/RG.2.2.16602.09921

[42] R. Parasuraman, T. Sheridan, and C. Wickens, "A model for types and levels of human interaction with automation," *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, vol. 30, no. 3, pp. 286–297, 2000.

[43] S. Raisch and S. Krakowski, "Artificial intelligence and management: The automation–augmentation paradox," *Academy of Management Review*, vol. 46, no. 1, pp. 192–210, 2021.

[44] S. Hong, M. Zhuge, J. Chen, X. Zheng, Y. Cheng, C. Zhang, J. Wang, Z. Wang, S. K. S. Yau, Z. Lin, L. Zhou, C. Ran, L. Xiao, C. Wu, and J. Schmidhuber, "Metagpt: Meta programming for a multi-agent collaborative framework," arXiv preprint, 2023.

[45] X. Tang, A. Zou, Z. Zhang, Z. Li, Y. Zhao, X. Zhang, A. Cohan, and M. Gerstein, "Medagents: Large language models as collaborators

for zero-shot medical reasoning," in *Findings of the Association for Computational Linguistics: ACL (Findings of ACL)*, 2024, pp. 599–621.

[46] J. Gottweis, W.-H. Weng, A. Daryin, T. Tu, P. Sirkovic, A. Myaskovsky, G. Glowaty, F. Weissenberger, A. Orlandi, D. Popovici, A. K. Palepu, K. Rong, R. Tanno, K. K. Saab, F. Zhang, J. A. Blum, A. W. Carroll, K. Kulkarni, N. Tomašev, D. Zverinski, I. Rendulic, E. Vedadi, F. Hasler, L. Rimanić, M. Boia, I. Budiselic, B. Feinstein, M. Bellaiche, T. Sheffer, J. Freyberg, J. Ratcliff, O. Bertolli, K. Chou, A. Hassidim, S. B. Gokturk, A. Vahdat, Y. Guan, V. Dhillon, E. D. Vaishnav, B. Lee, T. R. D. Costa, J. R. Penadés, G. Peltz, Y. Matias, J. Manyika, D. Hassabis, Y. Xu, P. Kohli, A. Pawlosky, A. Karthikesalingam, and V. Natarajan, "Accelerating scientific discovery with co-scientist," *Nature*, vol. 655, no. 8122, pp. 487–496, 2026.

[47] Z. Xi, W.-X. Chen, X. H. Guo, W. He, Y. Ding, B. Hong, M. Zhang, J. Wang, S. Jin, E. Zhou, R. Zheng, X. Fan, X. Wang, L. Xiong, Y. Z. Zhou, W. Wang, C. Jiang, Y. Zou, X. Liu, Z. Yin, S. Dou, R. Weng, W. Qin, Y. Zheng, X. Qiu, X. Huang, Q. Zhang, and T. Gui, "The rise and potential of large language model based agents: a survey," *Science China Information Sciences*, vol. 68, no. 2, 2025.

[48] R. Sapkota, K. I. Roumeliotis, and M. Karkee, "Ai agents vs. agentic ai: A conceptual taxonomy, applications and challenges," *Information Fusion*, vol. 126, p. 103599, 2025.

[49] A. Rah, "Agentic and multi-agent systems: A systematic review of tool use, benchmarks, and governance," Research Square preprint, 2026. [Online]. Available: https://doi.org/10.21203/rs.3.rs-9839978/v1

[50] ——, "Agentic artificial intelligence: A systematic review of architectures, benchmarks, and governance frameworks," Research Square preprint, 2026. [Online]. Available: https://doi.org/10.21203/rs.3.rs-9839527/v1

[51] K.-H. Huang, A. Prabhakar, O. Thorat, D. Agarwal, P. K. Choubey, Y. Mao, S. Savarese, C. Xiong, and C.-S. Wu, "Crmarena-pro: Holistic assessment of llm agents across diverse business scenarios and interactions," arXiv preprint, 2025.

[52] A. Drouin, M. Gasse, M. Caccia, I. Laradji, M. D. Verme, T. Marty, L. Boisvert, M. Thakkar, Q. Cappart, D. Vázquez, N. Chapados, and A. Lacoste, "Workarena: How capable are web agents at solving common knowledge work tasks?" arXiv preprint, 2024.

[53] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. V. Le, G. E. Hinton, and J. Dean, "Outrageously large neural networks: The sparsely-gated mixture-of-experts layer," arXiv preprint, 2017.

[54] Y. Moslem and J. D. Kelleher, "Dynamic model routing and cascading for efficient llm inference: A survey," arXiv preprint, 2026.

[55] A. Rah, "A digital-twin multi-agent framework for enterprise cognitive networking: ECNetBench and leakage-safe multi-seed evaluation," SSRN preprint, 2026. [Online]. Available: https://doi.org/10.2139/ssrn.7260325

[56] ——, "Agent-to-agent business ecosystems: Delegated commercial agency, governance, and a reference architecture for autonomous enterprise commerce," SSRN preprint, 2026. [Online]. Available: https://doi.org/10.2139/ssrn.7125338

[57] A. Rah and S. Hahues, "Retrieval-augmented generation (RAG), generative AI, and agentic AI governance: An integrated enterprise governance prioritization architecture," Research Square preprint, 2026. [Online]. Available: https://doi.org/10.21203/rs.3.rs-10012178/v1

[58] A. Rah, "Semantic tracing in LLM-based multi-agent systems using LangChain, LangGraph, and LangSmith for AI governance," Research Square preprint, 2026. [Online]. Available: https://doi.org/10.21203/rs.3.rs-10485157/v1

[59] T. Birkstedt, M. Minkkinen, A. Tandon, and M. Mäntymäki, "Ai governance: themes, knowledge gaps and future agendas," *Internet Research*, vol. 33, no. 7, pp. 133–167, 2023.

[60] A. Jobin, M. Ienca, and E. Vayena, "The global landscape of ai ethics guidelines," *Nature Machine Intelligence*, vol. 1, no. 9, pp. 389–399, 2019.

[61] A. Rah, "Enterprise generative AI: A systematic review of security risks and governance," Research Square preprint, 2026. [Online]. Available: https://doi.org/10.21203/rs.3.rs-9890887/v1

[62] J. Cobbe, M. S. A. Lee, and J. P. Singh, "Reviewable automated decision-making: A framework for accountable algorithmic systems," in *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2021, pp. 598–609.

[63] M. Suksi, "Administrative due process when using automated decision-making in public administration: some notes from a finnish perspective," *Artificial Intelligence and Law*, vol. 29, no. 1, pp. 87–110, 2020.

[64] F. Bignami, "Artificial intelligence accountability of public administration," *The American Journal of Comparative Law*, vol. 70, no. Suppl. 1, pp. i312–i346, 2022.

[65] B. J. Dietvorst, J. P. Simmons, and C. Massey, "Algorithm aversion: People erroneously avoid algorithms after seeing them err," *Journal of Experimental Psychology General*, vol. 144, no. 1, pp. 114–126, 2014.

[66] ——, "Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them," *Management Science*, vol. 64, no. 3, pp. 1155–1170, 2016.

[67] J. W. Burton, M.-K. Stein, and T. B. Jensen, "A systematic review of algorithm aversion in augmented decision making," *Journal of Behavioral Decision Making*, vol. 33, no. 2, pp. 220–239, 2019.

[68] H. Mahmud, A. N. Islam, S. I. Ahmed, and K. Smolander, "What influences algorithmic decision-making? a systematic literature review on algorithm aversion," *Technological Forecasting and Social Change*, vol. 175, p. 121390, 2021.

[69] M. J. Page, J. E. McKenzie, P. M. Bossuyt *et al.*, "The prisma 2020 statement: an updated guideline for reporting systematic reviews," *BMJ*, p. n71, 2021.

[70] M. L. Rethlefsen, S. Kirtley, S. Waffenschmidt *et al.*, "Prisma-s: an extension to the prisma statement for reporting literature searches in systematic reviews," *Systematic Reviews*, vol. 10, no. 1, 2021.

[71] A. Cooke, D. Smith, and A. Booth, "Beyond pico," *Qualitative Health Research*, vol. 22, no. 10, pp. 1435–1443, 2012.

[72] J. N. Ilonga and M. M. Ujakpa, "Leadership as an enabler of it governance in universities: Case study of namibian universities," *Namibia Journal of Managerial Sciences*, vol. 6, no. 3, pp. 112–127, 2025.

[73] V. M. d. Santos, L. F. Moura, and P. R. Zarelli, "Diagnostic model for it governance assessment in federal public universities," *Revista de Administração Contemporânea*, vol. 29, no. 3, 2025.

[74] L.-E. STANGACIU, "The current landscape of digital transformation in administrative processes of higher education: A systematic literature review," *Review of Economic and Business Studies*, vol. 18, no. 2, p. 215, 2025.

[75] R. Selvaratnam, K. Ames, and S. Leichtweis, "2024 governance of artificial intelligence and data in australasian higher education," *ASCILITE Publications*, 2024.

[76] Q. N. Hong, S. Fàbregues, G. Bartlett *et al.*, "The mixed methods appraisal tool (mmat) version 2018 for information professionals and researchers," *Education for Information*, vol. 34, no. 4, pp. 285–291, 2018.

[77] J. Thomas and A. Harden, "Methods for the thematic synthesis of qualitative research in systematic reviews," *BMC Medical Research Methodology*, vol. 8, no. 1, 2008.

[78] C. Carroll, A. Booth, J. Leaviss, and J. Rick, ""best fit" framework synthesis: refining the method," *BMC Medical Research Methodology*, vol. 13, no. 1, 2013.

[79] V. Braun and V. Clarke, "Using thematic analysis in psychology," *Qualitative Research in Psychology*, vol. 3, no. 2, pp. 77–101, 2006.

[80] M. Campbell, J. E. McKenzie, A. Sowden *et al.*, "Synthesis without meta-analysis (swim) in systematic reviews: reporting guideline," *BMJ*, p. l6890, 2020.

[81] C.-W. Liu, P. Huang, and H. C. Lucas, "Centralized it decision making and cybersecurity breaches: Evidence from u.s. higher education institutions," *Journal of Management Information Systems*, vol. 37, no. 3, pp. 758–787, 2020.

[82] D. L. Carraway, "Information technology governance maturity and technology innovation in higher education: Factors in effectiveness," NC Digital Online Collection of Knowledge and Scholarship (The University of North Carolina at Greensboro), 2015.

[83] C. K. Jim and H.-C. Chang, "The current state of data governance in higher education," *Proceedings of the Association for Information Science and Technology*, vol. 55, no. 1, pp. 198–206, 2018.

[84] H.-J. Kam, D. J. Kim, and W. He, "Should we wear a velvet glove to enforce information security policies in higher education?" *Behaviour & Information Technology*, vol. 41, no. 10, pp. 2259–2273, 2021.

[85] A. Latchu and S. Singh, "Navigating ict governance in south africa's provincial administrations – insights from the gito at the offices of the premier," *"System analysis and applied information science"*, no. 3, pp. 17–28, 2025.

[86] ——, "From governance failure to governance reform: practitioner-led pathways for strengthening ict governance in south africa's public sector," *"System analysis and applied information science"*, no. 3, pp. 11–22, 2026.

[87] M. Omari, F. Amoani-Arthur, and G. Essien, "Challenges in information system procurement in higher educational institutions in ghana," Journals & Books Hosting (International Knowledge Sharing Platform), 2016.

[88] I. S. Bianchi, R. D. Sousa, R. Pereira, and J. v. Hillegersberg, "Baseline mechanisms for it governance at universities," Repositório do ISCTE-IUL, 2017.

[89] A. R. Sengik, G. L. Lunardi, I. S. Bianchi, and G. C. Wiedenhöft, "Using design science research to propose an IT governance model for higher education institutions," *Education and Information Technologies*, 2022.

[90] J. R. Batmetan, "Agile IT Governance for Sustainable Digital Transformation in Universities," *International Journal of Information Technology and Education*, 2025.

[91] ——, "The Role of Application Quality and Integration in Agile IT Governance for Digital Transformation in Higher Education Institutions," *International Journal of Information Technology and Education*, 2026.

[92] T. Ngqondi and H. Mauwa, "Information technology governance model for a low resource institution with fragmented it portfolio," *South African Journal of Higher Education*, vol. 34, no. 3, 2020.

[93] E. Gorman, F. Kpedekpo, A. E. Botwe, and E. S. Adu-Opoku, "Digital transformation and administrative efficiency at the university of education wenneba, ghana: Barriers, enablers, and policy implications," *Journal of Strategic Management*, vol. 6, no. 4, pp. 27–43, 2026.

[94] A. Latchu and S. Singh, "State information technology agency (sita) and the paradox of centralized ict procurement: Reform or redundancy in south africa's public sector," DigitalCommons - Kennesaw State University (Kennesaw State University), 2026.

[95] C. L. Borgman and M. E. Martone, "Academic senate and it leadership: Shared governance in practice," eScholarship (California Digital Library), 2018.

[96] I. Scalabrin Bianchi, R. Dinis Sousa, and R. Pereira, "Information technology governance for higher education institutions: A multi-country study," *Informatics*, vol. 8, no. 2, p. 26, 2021.

[97] A. Chahid, S. Ahriz, K. El Guemmat, and K. Mansouri, "Developing digital capabilities through it governance: a pls-sem analysis in moroccan higher education institutions," *Bulletin of Electrical Engineering and Informatics*, vol. 14, no. 2, pp. 1543–1552, 2025.

[98] I. S. Bianchi and R. D. Sousa, "It governance maturity in higher education: A study in brazilian and portuguese universities," Journal of the Association for Information Systems, 2018.

[99] I. O. Adam, H. Mohammed, M.-O. Kevor, and Y. B. Osei, "Does governance translate to transformation? direct and mediated pathways to digital government maturity in ghana's public sector," Journal of the Association for Information Systems, 2026.

[100] M. López-Nores, F. de Arriba-Pérez, and S. García-Méndez, "Rethinking administrative software procurement in medium-sized universities: A socio-technical case study of embedded in-house development," *Systems*, vol. 14, no. 8, p. 992, 2026.

[101] Z. Sakyoud, A. Aaroud, and K. Akodadi, "Optimization of purchasing business process in moroccan public universities based on cobit and artificial intelligence techniques," *Kybernetes*, vol. 53, no. 5, pp. 1607–1635, 2023.

[102] M. Kaya, "Digital transformation in public universities: Bpmn-based process automation and enterprise architecture redesign," *Alphanumeric Journal*, vol. 14, no. 1, pp. 13–40, 2026.

[103] S. Y. Hua, "Procurement maturity and it failures in the public sector," *Transforming Government: People, Process and Policy*, vol. 16, no. 4, pp. 554–566, 2022.

[104] A. Gjaerum, "Digital transformation in local government – unpacking challenges in buyer-provider relations," *Transforming Government: People, Process and Policy*, vol. 20, no. 2, pp. 188–205, 2025.

[105] M. M. Young, J. B. Bullock, and J. D. Lecy, "Artificial discretion as a tool of governance: A framework for understanding the impact of artificial intelligence on public administration," *Perspectives on Public Management and Governance*, 2019.

[106] G. Prencipe, A. Tommasi, C. Zavattari, G. Tesi, L. Storchi, and K. Shahin, "Transforming public administration workflows with multi-agent ai: A human-in-the-loop knowledge orchestration framework," *Computers in Human Behavior Reports*, vol. 23, p. 101253, 2026.

[107] V. H. Fernández-Bedoya, A. G. Mejía-Osorio, and E. J. Rojas-Villanueva, "Decision-making and governance in public universities in the era of artificial intelligence: A systematic review," *Administrative Sciences*, vol. 16, no. 7, p. 334, 2026.

[108] S. Tianzi and M. J. D. Pastor, "Ai managerialism in higher education: Algorithmic governance, shared authority, and the ethics of institutional control," *American Journal of Education and Technology*, vol. 5, no. 3, pp. 131–136, 2026.

[109] C. N. Günal, "Sustainable digital governance of ai-based decision support systems in public administration: A socio-technical capacity framework," *Sustainability*, vol. 18, no. 16, p. 8439, 2026.

[110] M. K. Damarched, "Agentic ai modernization: Transforming institutional infrastructure through orchestrated multi-agent llm framework," *Journal of Computer Science and Technology Studies*, vol. 8, no. 4, pp. 01–24, 2026.

[111] J. Fattah-Weil, B. Bender, and N. Gronau, "Stakeholder-driven ERP implementation in German higher education: key challenges and adoption factors," *Information Systems and E-Business Management*, 2026.

[112] N. Zaki, A. R. Alneyadi, M. Alamoodi, S. Alshaali, F. Alkaabi, and A. Alhashmi, "An AI-enabled framework for reducing administrative bureaucracy in higher education while maintaining governance and compliance," *Discover Artificial Intelligence*, 2026.

[113] F. X. Valverde Alulema and F. Llorens, "Alignment of the portfolio of it projects with the it governance in spanish universities," *TECHNO REVIEW. International Technology, Science and Society Review /Revista Internacional De Tecnología, Ciencia Y Sociedad*, vol. 10, no. 2, pp. 111–131, 2021.

[114] P. K. Banua, "Digital governance and state power in higher education: A political analysis of odisha's higher education information management system (hims)," *Integrated Journal for Research in Arts and Humanities*, vol. 6, no. 5, pp. 245–253, 2026.

[115] M. Abdulbari, "AI Governance in Higher Education: Legal Implications for University Autonomy," *European journal of law and technology*, 2026.

[116] M. Patalon, "Situated isomorphism: institutional logics and the interpretive dynamics of municipal digital transformation," *Journal of Organizational Change Management*, vol. 39, no. 8, pp. 39–55, 2026.

[117] V.-H. C. X.-L. Pham, "Rethinking Technology Acceptance in Automation Contexts: Evidence from Robotic Process Automation Adoption in Vietnamese Higher Education Institutions," *Journal of Intelligent Decision Making and Information Science*, 2026.

[118] Y. M. N. Almutairi, "The Impact of Digital Transformation on the Saudi Higher Education Administration," *International Journal of Instruction*, 2026.

[119] A. Joodala, "Legacy System Decomposition Techniques for Public Sector Cloud Migrations," *International Journal of Innovative Research in Engineering & Multidisciplinary Physical Sciences*, 2025.

[120] Y. Soegoto, H. Prabowo, H. Soeparno, and H. L. H. S. Warnars, "Fragmented Systems to Integrated Digital Transformation: An Enterprise Architecture Model for Indonesian Private Higher Education," *ASEAN Journal of Educational Research and Technology*, 2026.

[121] A. Almoqry and N. T. Alsohybe, "Enterprise Architecture in Higher Education for Digital Transformation: An Integrative Review for Fragile Contexts with Application to Yemen," *F1000Research*, 2025.

[122] Taryana, "Organizational barriers to data-driven decision making in educational management," *Journal of Economics, Entrepreneurship, Management Business and Accounting*, vol. 4, no. 3, pp. 457–477, 2026.

[123] H. Ronoh and A. Isiaho, "SECURE MODEL FOR ERP-CLOUD INTEGRATION FOR SUSTAINABLE DIGITAL TRANSFORMATION IN KENYAN UNIVERSITIES," *IJARCCE*, 2025.

[124] G. L. D. Santos, M. A. Oliveira, and A. Bueno, "Digital transformation readiness in university operations: A product-service system approach to readiness-as-a-service," *IEEE Access*, vol. 14, pp. 60 190–60 204, 2026.

[125] T. Nguyen, "Digital Transformation in a Vietnamese Public University: Beyond Software Acquisition," *International Journal of Advanced Multidisciplinary Research and Studies*, 2026.

[126] I. Kumwenda, J. Likando, M. Malama, H. Siame, and D. Mainde, "Contextual Barriers to AI Governance Adaptation in Zambian Private Higher Education," *International Journal of Research and Innovation in Social Science*, 2026.

[127] K. Asmar and L. BELFAKIR, "Strategic capability and portfolio governance in higher education digital transformation: A systematic literature review," International Journal Of Applied Management And Economics, 2026.

[128] D. S. G. Dr. Sam Ghosh, "University it project alignment with it governance framework," *International Journal of Business and Management Invention*, vol. 14, no. 8, pp. 53–61, 2025.

[129] B. Attard-Frost, A. Brandusescu, and K. Lyons, "The governance of artificial intelligence in canada: Findings and opportunities from a review of 84 ai governance initiatives," *Government Information Quarterly*, vol. 41, no. 2, p. 101929, 2024.

[130] M. G. Özsavaş and E. E. Özsavaş, "Ai governance capacity in public administration: A comparative analysis of the eu and türkiye," *International Journal of Public Administration in the Digital Age*, vol. 13, no. 1, 2026.

[131] E. Şen, D. Vaněček, and M. Adnan, "Balancing regulation and innovation: the need for agile ai governance in higher education – a cross-country study," *Studies in Higher Education*, 2026.

[132] A. N. Dragomir, I. Bulea, and L. I. Tarnu, "Human-centric ai governance in the european union. accountability, fundamental rights, and institutional resilience in public administration," *Frontiers in Political Science*, vol. 8, 2026.

[133] T. Pariwat, P. Chantaramongkol, P. Ruangkulsap, J. Chitiyaphol, T. Jantakoon, and W. Unarat, "Development of an AI-Driven Educational Management Information System Architecture for Data-Driven Decision-Making in Higher Education," *Higher Education Studies*, 2026.

[134] I. García-López and J. Nájera-Ochoa, "Algorithmic Governance and Strategic Value of Open Large Language Models in University Administrative Decision-Making," *Revista Universidad y Empresa*, 2026.

[135] R. M. Aguilar, S. Alayón, J. M. Torres, and C. A. Martín, "Data-centric AI governance for responsible organizational value: evidence from a European public administration," *AI & Society*, 2026.

[136] N. Pandey and D. Davies, "Higher education governance under ai and geopolitical disruption: a socio-technical framework for administrative artificial intelligence," *Frontiers in Education*, vol. 11, 2026.

[137] Z. Khajoueinejad, G. Shams, and F. Seraji, "Empowering Leaders for Human-Centered Digital Transformation: A Strategic Framework for Establishing Digital Universities," *Journal of Higher Education Policy And Leadership Studies*, 2025.

[138] V. Vydia, R. Gernowo, and A. Hendrawan, "A Knowledge Management Integration Model within Enterprise Architecture to Improve the Competitiveness of Higher Education Institutions," *JINAV Journal of Information and Visualization*, 2026.

[139] I. R. A. Shamsi, "Empowering Higher Education Through Digital Transformation and Strategic Planning for Academic Advancement," *Emerging Science Journal*, 2025.

[140] S. Khan, M. Saif, and N. Sharma, "Sustainable Development Practices for Adoption of Robotic Process Automation (RPA) at the Universities," *Journal of Interdisciplinary Knowledge*, 2025.

[141] A. Rah and D. W. Johnson, "Risk-aware AI procurement in the U.S. public sector: Stochastic optimization and technology governance," Research Square preprint, 2026. [Online]. Available: https://doi.org/10.21203/rs.3.rs-11679788/v1

[142] B. K. AlNuaimi, S. Kumar Singh, S. Ren, P. Budhwar, and D. Vorobyev, "Mastering digital transformation: The nexus between leadership, agility, and digital strategy," *Journal of Business Research*, vol. 145, pp. 636–648, 2022.

[143] R. K. R. A. H. K. A. J. S. M. A. A. F. Mudhafar, "Implementing agile principles in the public sector: An analysis of barriers and the role of leadership," *Acta Innovations*, vol. 62, no. 1s, 2026.

[144] D. Baladi and B. Guennoun, "Governance of inclusive decision-making in AI adoption: Participation and institutional readiness in Moroccan universities," *Journal of Interdisciplinary Studies in Education*, 2026.

[145] L. B. Carvalho, F. Buchbinder, D. A. B. da Hora, V. B. Ferreira, and D. M. de Carvalho, "Social participation as an ethical criterion for AI governance in public administration," *AI and Ethics*, 2026.

[146] J. R. Batmetan, "Sustainable Digital Transformation in Indonesian Higher Education: The Role of Service Planning and Architecture Strategy within an Agile IT Governance Approach," *TADBIR Jurnal Studi Manajemen Pendidikan*, 2026.

[147] M. Karakaya, "A design science framework for ethical data-driven decision-making in higher education," *Discover Artificial Intelligence*, 2026.

[148] E. Molchanova, "Core eu principles for digital transformation in higher education institutions: Governance, trust, and interoperability," *Pedagogy and Education Management Review*, no. 4(22), pp. 4–16, 2025.

[149] K. TN and N. Dhanraj, "Digital Transformation in Higher Education: An Empirical Study of AI-Integrated ERP Systems in Enhancing Decision-Making and Institutional Efficiency," *International Journal of Economic Practices and Theories*, 2026.

[150] M. K. Damarched, "Applying llms to legacy system modernization in higher education it: Leveraging large language models beyond chatbots to modernize core student and administrative systems in universities—a suggestive review study," *International Journal of Innovative Science and Research Technology*, p. 3043, 2026.

[151] A. Rah, "AI and quantum trading agents: Cybersecurity, governance, and regulatory challenges in financial decision-making," Research Square preprint, 2026. [Online]. Available: https://doi.org/10.21203/rs.3.rs-10253799/v1

[152] EDUCAUSE, "EDUCAUSE QuickPoll results: AI-related procurement," https://er.educause.edu/articles/2025/5/educause-quickpoll-results-ai-related-procurement, 2025, accessed: Nov. 15, 2026.

[153] Texas Quality Assurance Team (Comptroller of Public Accounts, Department of Information Resources, Legislative Budget Board, and State Auditor's Office), "2025 Quality Assurance Team annual report: Overview of major information resources projects," Legislative Budget Board ID 9365, https://www.lbb.texas.gov/Documents/Publications/Policy_Report/2025_QAT_Annual_Report.pdf, 2025, 59 major IT projects ($1.9B) in FY2025; accessed: Nov. 15, 2026.

[154] B. van Dongen, "BPI challenge 2020: Request for payment and travel permit event logs (Eindhoven university of technology)," 4TU.ResearchData, CC BY-NC 4.0, https://doi.org/10.4121/uuid:895b26fb-6f25-46eb-9e48-0dca26fcd030 and https://doi.org/10.4121/uuid:ea03d361-a7cd-4f5e-83d8-5fbdf0362550, 2020, mD5 2eb4dd20e70b8de4e32cc3c239bde7f2 and b6e9ff00d946f6ad4c91eb6fb550aee4.

[155] University of Washington Information Technology, "UW enterprise IT projects: Project portfolio executive summary (12/31/21)," University of Washington Office of the CIO / Technology Services Board, https://it.uw.edu/wp-content/uploads/2022/07/UW-Project-Portfolio-Exec-Summary-Dec-2021-FINAL-26jan22.pdf, 2022, 11 university enterprise IT projects ($368.5M) across Oversight Levels 1–3; accessed: Nov. 15, 2026.

[156] S. Alon-Barkat and M. Busuioc, "Human–ai interactions in public sector decision making: "automation bias" and "selective adherence" to algorithmic advice," *Journal of Public Administration Research and Theory*, vol. 33, no. 1, pp. 153–169, Feb. 2022. [Online]. Available: http://dx.doi.org/10.1093/jopart/muac007

[157] R. F. Frogeri, D. J. Pardini, and G. R. Cunha, "Information technology governance in a higher education institution," *International Journal of Human Capital and Information Technology Professionals*, vol. 11, no. 1, pp. 31–46, 2020.

[158] U.S. Department of Education, "34 CFR § 99.31, under what conditions is prior consent not required to disclose information?" https://www.ecfr.gov/current/title-34/subtitle-A/part-99/subpart-D/section-99.31, accessed: Nov. 15, 2026.

[159] U.S. Department of Health and Human Services, "Covered entities and business associates," https://www.hhs.gov/hipaa/for-professionals/covered-entities/index.html, accessed: Nov. 15, 2026.

[160] U.S. Department of Justice, Civil Rights Division, "Fact sheet: New rule on the accessibility of web content and mobile apps provided by state and local governments (ADA title II)," https://www.ada.gov/resources/2024-03-08-web-rule/, accessed: Nov. 15, 2026; reflects the interim final rule of Apr. 20, 2026.

[161] U.S. Government, "2 CFR § 200.318, general procurement standards," https://www.law.cornell.edu/cfr/text/2/200.318, accessed: Nov. 15, 2026.

[162] General Services Administration, Department of Defense, and NASA, "Federal acquisition regulation, 1.104 applicability," https://www.acquisition.gov/far/1.104, fAC 2026-01; accessed: Nov. 15, 2026.

[163] Executive Office of the President, Office of Management and Budget, "Accelerating federal use of AI through innovation, governance, and public trust," OMB, Tech. Rep. M-25-21, 2025.

[164] S. D. Young, "Advancing governance, innovation, and risk management for agency use of artificial intelligence," Executive Office of the President, Office of Management and Budget, Tech. Rep. M-24-10, 2024, rescinded and replaced by M-25-21.

[165] Federal Trade Commission, "16 CFR § 314.1, purpose and scope (standards for safeguarding customer information)," https://www.law.cornell.edu/cfr/text/16/314.1, accessed: Nov. 15, 2026.

[166] Federal Student Aid, U.S. Department of Education, "Updates to the Gramm-Leach-Bliley Act cybersecurity requirements (electronic announcement general-23-09)," https://fsapartners.ed.gov/knowledge-center/library/electronic-announcements/2023-02-09/updates-gramm-leach-bliley-act-cybersecurity-requirements, 2023, accessed: Nov. 15, 2026.

[167] ——, "Service provider relationships for GLBA compliance," https://fsapartners.ed.gov/knowledge-center/library/electronic-announcements/2024-04-24/service-provider-relationships-glba-compliance, 2024, accessed: Nov. 15, 2026.

[168] Texas Department of Information Resources, "1 tex. admin. code ch. 202, subch. c, information security standards for institutions of higher education, §§ 202.70–202.77," https://www.law.cornell.edu/regulations/texas/1-Tex-Admin-Code-SS-202-74, 1 TAC § 202.74; transfer to ch. 477 noted; accessed: Nov. 15, 2026.

[169] Texas Legislature, "Tex. bus. & com. code § 552.001 (texas responsible artificial intelligence governance act, H.B. 149, eff. jan. 1, 2026)," https:

//statutes.capitol.texas.gov/GetStatute.aspx?Code=BC&Value=552, accessed: Nov. 15, 2026.
[170] S. D. Haes and W. V. Grembergen, "It governance structures, processes and relational mechanisms: Achieving it/business alignment in a major belgian financial group," in *Hawaii International Conference on System Sciences (HICSS)*, 2005, p. 237b.
[171] A. Schwarz and R. Hirschheim, "An extended platform logic perspective of it governance: managing perceptions and activities of it," *The Journal of Strategic Information Systems*, vol. 12, no. 2, pp. 129–166, 2003.
[172] B. van Dongen, "BPI challenge 2019," 4TU.ResearchData dataset, CC BY 4.0, https://doi.org/10.4121/uuid:d06aff4b-79f0-45e6-8ec8-e19730c248f1, 2019, bPI_Challenge_2019.xes; MD5 4eb909242351193a61e1c15b9c3cc814.
[173] Z. Hu and L. J. Hong, "Kullback-Leibler divergence constrained distributionally robust optimization," Optimization Online preprint, 2013.
[174] D. Kuhn, S. Shafiee, and W. Wiesemann, "Distributionally robust optimization," *Acta Numerica*, vol. 34, pp. 579–804, Jul. 2025. [Online]. Available: http://dx.doi.org/10.1017/s0962492924000084
[175] H. Rahimian and S. Mehrotra, "Frameworks and results in distributionally robust optimization," *Open Journal of Mathematical Optimization*, vol. 3, pp. 1–85, Jul. 2022. [Online]. Available: http://dx.doi.org/10.5802/ojmo.15
[176] F. Lin, X. Fang, and Z. Gao, "Distributionally robust optimization: A review on theory and applications," *Numerical Algebra, Control & Optimization*, vol. 12, no. 1, p. 159, 2022. [Online]. Available: http://dx.doi.org/10.3934/naco.2021057
[177] R. Lahdelma and P. Salminen, "Smaa-2: Stochastic multicriteria acceptability analysis for group decision making," *Operations Research*, vol. 49, no. 3, pp. 444–454, Jun. 2001. [Online]. Available: http://dx.doi.org/10.1287/opre.49.3.444.11220
[178] R. Lahdelma, K. Miettinen, and P. Salminen, "Ordinal criteria in stochastic multicriteria acceptability analysis (smaa)," *European Journal of Operational Research*, vol. 147, no. 1, pp. 117–127, May 2003. [Online]. Available: http://dx.doi.org/10.1016/s0377-2217(02)00267-9
[179] Enterprise Information Services, State of Oregon, "Enterprise information services IT project portfolio (stage-gate oversight dataset)," Oregon Open Data Portal, U.S. Government Works, https://data.oregon.gov/d/hjrz-mzrm, 2026, snapshot Nov. 3, 2026 (115 active/oversighted projects); accessed: Nov. 15, 2026.
[180] City of Mesa Department of Information Technology, "Information technology — project schedule and budget variance dataset," City of Mesa Open Data Portal, https://citydata.mesaaz.gov/d/24pf-5hpg, 2026, 1,645 quarterly records across 256 IT projects; accessed: Nov. 15, 2026.
[181] E. Riihimäki and S. Pekkola, "Public buyer's concerns influencing the early phases of information system acquisition," *Government Information Quarterly*, vol. 38, no. 4, p. 101595, 2021.
[182] G. Chiandussi, M. Codegone, S. Ferrero, and F. Varesio, "Comparison of multi-objective optimization methodologies for engineering applications," *Computers & Mathematics with Applications*, vol. 63, no. 5, pp. 912–942, Mar. 2012. [Online]. Available: http://dx.doi.org/10.1016/j.camwa.2011.11.057
[183] Z. S. Hosseini, S. D. Flapper, and M. Pirayesh, "Sustainable supplier selection and order allocation under demand, supplier availability and supplier grading uncertainties," *Computers & Industrial Engineering*, vol. 165, p. 107811, Mar. 2022. [Online]. Available: http://dx.doi.org/10.1016/j.cie.2021.107811
[184] M. A. Naqvi and S. H. Amin, "Supplier selection and order allocation: a literature review," *Journal of Data, Information and Management*, vol. 3, no. 2, pp. 125–139, May 2021. [Online]. Available: http://dx.doi.org/10.1007/s42488-021-00049-z
[185] D. A. Dung, N. T. Toan, N. D. Minh, and N. Q. Anh, "Reimagining quality: Artificial intelligence, governance and the politics of data in higher education," *Higher Education for the Future*, vol. 13, no. 1, pp. 50–73, 2025.
[186] Y. R. E. Dongallo, "Artificial Intelligence-Assisted Strategic Decision-Making for Sustainable University Governance," *Engineering and Technology Journal*, 2026.
[187] W. Zhang and O. Rashdi Syed, "Global bibliometric mapping of faculty resistance to change in higher education: Uncovering the psychological costs of university governance," *International Journal of Academic Research in Business and Social Sciences*, vol. 16, no. 3, 2026.
[188] J. Hong, ""run at most once" reform in higher education: Change from informatization to digitalization of administration in university x," *The Educational Review, USA*, vol. 7, no. 11, pp. 1821–1826, 2023.
[189] I. S. Bianchi, R. D. Sousa, R. Pereira, and S. De Haes, "The influence of culture in it governance implementation," *International Journal of Human Capital and Information Technology Professionals*, vol. 10, no. 4, pp. 55–68, 2019.

APPENDIX A
COMPARISON OF THEORETICAL LENSES

TABLE XXVI
THEORETICAL LENSES: WHAT THEY EXPLAIN, USE IN INCLUDED STUDIES, AND LIMITS FOR THIS DOMAIN

Lens	What it explains	Use in the included studies	Limit for approval-chain research
Institutional theory [20]	Convergence of structures under coercive, mimetic, and normative pressure; symbolic compliance.	6 full-text studies (e.g., [4], [7]).	Explains persistence, not whether a control is necessary.
Path dependence and structural inertia [21], [22]	Self-reinforcing sequences; slow change relative to environments.	Implicit in legacy/bureaucracy analyses (e.g., [102], [119]).	Requires longitudinal event-log data to separate protective vs. deadweight delay.
Organized anarchy and loose coupling [23], [24]	Irregular problem–solution matching and weak links between units in universities.	Not primary in the included studies; consistent with decoupling findings [99].	Descriptive; offers few design levers.
Technology acceptance (TAM, UTAUT) [25], [26]	Individual intention and use.	Not primary in any included study.	Unit of analysis is the user, not the decision process.
Dynamic capabilities [27] and digital transformation research [28]	Organizational sensing, seizing, and reconfiguring.	Related capacity themes (e.g., [97]).	Weak on legal and safeguard constraints.
Public value governance [30]	Legitimacy and accountability alongside efficiency.	Related use [143].	Few operational measures.
Sociotechnical systems [31]	Joint design of technical and social arrangements.	About 4 studies (e.g., [136], [185]).	Qualitative; limited predictive use.
IT governance and COBIT [32]	Decision rights, controls, maturity.	About 20 studies; 13 use COBIT explicitly.	Maturity measurement without outcome linkage.

APPENDIX B
REPRODUCIBILITY AND EXECUTABLE DELIVERABLES

The project archive includes: (i) the 14-query uncapped search log (`search_log.csv`), screening ledger and reconciled PRISMA count files (`screening_ledger_11265.csv`, `exclusion_reasons.csv`, `prisma_counts.json`), PRISMA 2020 / PRISMA-S checklists (`prisma_checklist.md`, `prisma_s.md`), full-text exclusion and ILL logs (`fulltext_exclusions.csv`, `reports_not_retrieved_ill_list.csv`, `66_studies_access_links.pdf`, `second_pass_retrieval_report.md`), full-text extraction (`evidence_extraction.csv`), item-by-item MMAT 2018 / JBI appraisal (`formal_mmat_jbi_appraisal.csv`), and the search-strategy note (`search_strategy.md`); (ii) checksum-verified public empirical governance datasets under `RA_GO/data_empirical/` with process-mining and audit scripts under `RA_GO/scripts/` (including `bpi2019_process_mining.py`); and (iii) RA-GO v2 numerical-experiment, fairness, OOD, and exact 2^3 Shapley scripts under `RA_GO/`, and the independent Evaluation D suite (`RA_GO/eval_suite/`). Evaluation C frozen configuration identifiers (model, seed, prompt and view SHA-256 prefixes) are recorded in Appendix D; raw Evaluation C JSONL run files are not included in the archive.

APPENDIX C
FULL-TEXT-VERIFIED INCLUDED STUDIES

Tables XXVII and XXVIII list all 69 full-text-verified English studies with their setting/sample, full-text study design, verified finding, and formal MMAT 2018 or JBI checklist score (criteria met / total and quality tier).

APPENDIX D
REPRODUCIBILITY OF EVALUATIONS C AND D

This appendix records the frozen configurations and count tables needed to reproduce Evaluations C and D. Evaluation C uses labels distinct from Evaluation B's rule-derived D-sample ($200,689$ closed cases). Evaluation D is independent of Evaluation A.

A. Evaluation C: Frozen Configuration

A locally hosted 8-billion-parameter instruction-tuned model (Qwen3 8B) was run once on the test partition with temperature 0, seed 7, context window $4,096$ tokens. Prompt SHA-256 prefix `4e8de0563ee327aa`; frozen record-view SHA-256 prefix `6d511c4e14caa233`. Purchasing documents were split by document start time (test: last 40%, cutoff 2018-07-26). Prompts contained the shared rule text and discrepancy-code list but not stratum, injection type, or gold labels. Events were re-sorted by timestamp after injection.

B. Evaluation C: Detection Counts and Strata

Positives $= 78$ (30 real D + 48 injected S); negatives $= 60$ clean-by-rule. LLM-RO: TP/FP/FN/TN $= 78/59/0/1$. LLM-SA: $55/13/23/47$. LLM-CM: $55/12/23/48$. LLM-ME: $73/42/5/18$. Wilson 95% CIs: LLM-SA recall 60–79%, FPR 13–34%; LLM-ME recall 86–97%, FPR 57–80%.

Recall on real D ($n = 30$): LLM-RO/SA/CM $30/30$; LLM-ME $29/30$ (96.7%). Injected rule-visible I1/I2/I4 ($n = 24$): SA/CM $17/24$; ME $23/24$. Injected types not in the documented rules I3/I5/I6 ($n = 24$): SA/CM $8/24$; ME $21/24$. Per

injected type (8 each): I1 and I2 were $8/8$ for every LLM arm; I4 was $1/8$ (SA/CM) vs. $7/8$ (ME); I3 $4/8$ vs. $8/8$; I5 $2/8$ vs. $6/8$; I6 $2/8$ vs. $7/8$. Type-correct recall (flagged and a correct code named) was 58–67% for all LLM arms.

Exact McNemar tests (Holm over 18 pair×scope tests): LLM-SA vs. LLM-ME overall correctness 31 vs. 20 ($p = 0.1608$, Holm 0.64); positives-only 1 vs. 19 (Holm 0.0004); negatives-only 30 vs. 1 (Holm < 0.001). LLM-SA vs. LLM-CM overall: 0 vs. 1 ($p = 1$).

Mean per-case tokens and sequential latency: LLM-RO $1,072$ / 8.7 s; LLM-SA $1,228$ / 14.1 s; LLM-CM $4,109$ / 40.4 s (85% of LLM-ME tokens); LLM-ME $4,828$ / 41.9 s. Automatic JSON/governance checks: LLM-CM had one invalid discrepancy-code string; other LLM arms 0.

Confusion-matrix counts in Fig. 10 match Table XVI; raw Evaluation C JSONL run files are not included in the archive (configuration identifiers above).

C. Evaluation D: Gold Lock and Held-Out Faults

Gold SHA-256 prefix `861eb45130705752`; freeze assigned `gold_`$\star$ from complete facts before any arm was scored. Held-out halt faults (disjoint from the production GOV registry `invoice_before_gr_when_after_required` and `cleared_before_gr`): quantity mismatch (FIN invoice vs. PROC GR), spend vs. sealed-bid threshold (FIN vs. LEGAL), sanctions with payment pending (PROC vs. FIN), and residency with unapproved region (SEC vs. TECH). Optional-token dropout 0.1 over seeds 0–4 left S and M_{join} identical to each other; accuracy fell together (range 85.3–97.1%). Five-class and human-review matrices are Fig. 11.

APPENDIX E
RA-GO V2 SCENARIO ASSUMPTIONS

Table XXIX lists the 20 constructed scenarios. Alternatives, costs, exposures, and delays per alternative are in the supplementary scenario dataset (`RA_GO/data/scenarios.csv`); every numeric value is model-generated, and the evidence anchors in that file justify qualitative settings only.

TABLE XXVII
INCLUDED STUDIES (PART 1 OF 2; $N = 69$): SETTING, DESIGN, FINDING, AND MMAT 2018 / JBI CRITERIA MET

Ref.	Yr	Setting / sample	Design (full text)	Full-text-verified focus / finding	Items (%)
[115]	2026	UK / European higher education	Doctrinal & multi-layer legal-governance analysis	Analyzes interplay of statutory data protection (e.g., UK GDPR automated decision rules), soft law, and university autonomy; requires transparency and contestability for consequential AI decisions	6/6 (100%)
[135]	2026	Europe; 1 public administration agency (legislative AI system)	Qualitative empirical case study	Case study of an automated legislative-monitoring AI system demonstrates transition from ad-hoc model experimentation to data governance and MLOps with continuous human oversight	4/5 (80%)
[7]	2026	Saudi Arabia; 24 interviews	Qualitative grounded theory	Four response modes to coerced DT incl. surface compliance and active resistance	4/5 (80%)
[118]	2026	Saudi Arabia; 6 public & private universities (N=930 admin/IT staff)	Quantitative cross-sectional survey (TOE)	Across 930 administrative and IT staff in six universities, digital transformation positively relates to administrative efficiency, data management, and decision processes, constrained by legacy workflows and skill gaps	3/5 (60%)
[143]	2026	Public sector civil-service agencies; N=342	Mixed methods (PLS-SEM)	Authors report institutional rigidity, bureaucratic inertia and legacy systems negatively related to public value creation, with agile transformational leadership moderating the first two	2/5 (40%)
[127]	2026	Global; 16 articles	Systematic review (PRISMA-oriented)	DT often pursued at project level without portfolio governance; alignment and decision-making challenges	8/11 (73%)
[144]	2026	Morocco / global higher education (37 primary studies)	Systematic literature review (PRISMA)	PRISMA review of 37 empirical AI-governance studies in HE highlights participatory decision-making gaps and uneven institutional readiness	9/11 (82%)
[114]	2026	India; one state system	Policy and case analysis	State HE information system centralizes administration and data capability; raises accountability and privacy questions	3/5 (60%)
[91]	2026	Indonesia; Manado State University & HEIs	Mixed methods (PLS-SEM + case study)	Application interoperability, data consistency, and reliability mediate agile IT governance effectiveness in university digital transformation	5/5 (100%)
[146]	2026	Indonesia; 1 public university (N=12 stakeholder interviews + IT master plan)	Qualitative single-case study (SWOT, McFarlan grid)	Interviews with 12 leaders/IT managers and document analysis identify fragmented service planning and architecture silos; maps portfolio via McFarlan Strategic Grid	4/5 (80%)
[145]	2026	Global / public administration	Systematic literature review	Systematic review argues social participation functions as an ethical validation criterion for public-administration AI tools to strengthen legitimacy and accountability	8/11 (73%)
[150]	2026	Higher education IT	Suggestive narrative review	Full text (10,608 words) proposes LLM-assisted university legacy modernization architecture; contains secondary survey citations and illustrative metrics without reproducible primary institutional experiment	4/6 (67%)
[110]	2026	Higher education IT	Technical architecture proposal	Full text (10,359 words) proposes orchestrated multi-agent LLM legacy modernization framework; reported performance figures are architectural projections, not independently verified university trials	3/6 (50%)
[186]	2026	Philippines; university administrators	Quantitative descriptive-correlational survey	Administrator survey reports perceived alignment between AI-assisted strategic analytics and governance planning; self-report cross-sectional design without objective outcome metrics	3/5 (60%)
[132]	2026	EU public administration	Policy and practice review	Accountability must follow the chain of influence; protection uneven where audit capacity and procurement independence differ	6/6 (100%)
[111]	2026	Germany; 1 in-depth public university case + survey of 9 HEIs	Mixed methods (in-depth case + 9-HEI survey)	Identifies five determinants of German university ERP modernization: team resourcing, structured change management, process-ERP alignment vs customization, training, and sustained leadership	4/5 (80%)
[134]	2026	Global higher education (183 Scopus/WoS articles, 2019-2025)	Systematic literature review (PRISMA)	PRISMA review of 183 Scopus/WoS studies finds open LLMs in universities are used mainly for operational automation while strategic/governance mechanisms lag behind digital-skill and regulatory gaps	8/11 (73%)
[93]	2026	Ghana; 1 university, 40 interviews	Qualitative single case	Difficulties driven less by resource shortage than by strategy-operations misalignment, top-down communication, workarounds, ineffective training	4/5 (80%)
[6]	2026	Germany; 4 staff interviews	Qualitative case study	Resistance partly driven by speed and opacity of change and fear of the unknown	4/5 (80%)
[147]	2026	Higher education institutions	Design science research (ethical DDDM + TOPSIS)	Develops a 4-layer ethical data-driven decision-making architecture integrating SIS/HR/LMS repositories with algorithmic decision support (TOPSIS) and governance guardrails	6/6 (100%)
[102]	2026	Public universities (country not stated)	Process analysis (100 processes, BPMN) with quantitative comparison	Full text models 100 public-university administrative processes in BPMN (referencing TOGAF); documents multi-stage approval bottlenecks and proposes threshold-based approvals, e-signature, and RPA	2/5 (40%)
[126]	2026	Zambia; 5 private HEIs (N=100 governance staff survey + 12 leader interviews)	Sequential exploratory mixed methods	Mixed-methods study (100 governance staff, 12 leaders across 5 HEIs) finds AI governance constrained simultaneously from above (finance, infrastructure, regulatory gaps) and within (workflow habits)	5/5 (100%)
[72]	2025	Namibia; 3 universities; 9 interviews + 20 staff survey	Mixed methods (convergent parallel)	Dedicated leadership, digital proficiency, and collaborative networks enable IT governance; Cronbach $\alpha = 0.87$; 92% response rate	4/5 (80%)
[86]	2026	South Africa; 55 interviews	Qualitative interviews	Full text (55 interviews, South Africa) explicitly maps COBIT, ISO/IEC 38500, TOGAF, and ITIL; identifies five practitioner reform pathways including procurement streamlining and leadership accountability	5/5 (100%)
[125]	2026	Vietnam; 1 public university (Hanoi Univ. of Natural Resources, EFA/regression)	Quantitative single-institution survey (TOE + RBV)	Empirical survey at a Vietnamese public university shows digital transformation requires organizational, governance, and resource alignment beyond mere software acquisition	3/5 (60%)
[136]	2026	Global	Conceptual socio-technical framework	Administrative AI framework of drivers and governance conditions affecting resilience, equity, participation	6/6 (100%)
[133]	2026	Thailand; higher education (PRISMA synthesis + 7-expert CVI evaluation)	Design science research with expert content-validity testing	Develops AI-EMIS architecture; expert CVI evaluation passed 12/14 items while flagging AI explainability/traceability and feasibility under university constraints (I-CVI=0.71) as hardest	6/6 (100%)
[117]	2026	Vietnam; higher education institutions	Quantitative survey (PLS-SEM)	PLS-SEM study of administrative RPA adoption in Vietnamese HEIs highlights institutional readiness, rule-based process standardization, and heterogeneous infrastructure constraints	3/5 (60%)
[15]	2026	ASEAN; Scopus + gray	Scoping review (PRISMA-ScR)	More AI policies than empirical evaluations of them; policy-evidence mismatch framed as a governance-capacity problem	10/11 (91%)
[4]	2026	Global; 50 articles	Systematic review	Hierarchy remains important for legal certainty and accountability; modern governance combines stability with flexibility	5/11 (46%)

TABLE XXVIII
INCLUDED STUDIES (PART 2 OF 2; $N = 69$): SETTING, DESIGN, FINDING, AND MMAT 2018 / JBI CRITERIA MET

Ref.	Yr	Setting / sample	Design (full text)	Full-text-verified focus / finding	Items (%)
[122]	2026	Indonesia; N=179; mixed education institutions	Quantitative survey (SEM-PLS)	Data quality, organizational resistance, analytical capability deficiency and system fragmentation affect weak data-driven decisions; sample includes schools	4/5 (80%)
[108]	2026	UK, USA, Australia cases	Critical narrative review + 3 cases	AI can improve efficiency but risks opacity, surveillance, stakeholder exclusion and centralization of managerial authority	4/5 (80%)
[138]	2026	Higher education institutions (Indonesia)	Conceptual / Enterprise Architecture model (TOGAF, COBIT)	Integrates knowledge management into TOGAF/COBIT enterprise architecture to address early-stage university digital failure caused by strategy-technical misalignment	6/6 (100%)
[120]	2026	Indonesia; private higher education institutions	Design science research (EADT model)	DSR study documents lack of internal policy, budget limits, legacy/third-party fragmentation, security/privacy risks, weak interdepartmental coordination, and resistance to change	6/6 (100%)
[112]	2026	UAE; 1 public research university (UAEU, 5 HR workflows, 9-18 steps)	Multi-workflow design & comparative case evaluation	Evaluates 5-phase AI workflow redesign across 5 university HR processes (reporting 59.6% analytically modeled time reduction); relies on autonomous Step-3 execution and analytical pre/post modeling rather than controlled trial	6/6 (100%)
[187]	2026	Global; 48 Scopus records	Bibliometric mapping	Accountability and audit-oriented reforms linked in the literature to faculty resistance and psychological costs	8/11 (73%)
[121]	2025	Global / fragile-context HEIs with Yemen application (30 studies, 2010-2025)	Systematic integrative review (PRISMA; F1000 peer-reviewed)	PRISMA integrative review of 30 EA studies maps TOGAF, COBIT, and ISO/IEC 38500 governance mechanisms for resource-constrained universities facing fragmented legacy silos	10/11 (91%)
[90]	2025	Indonesia; multiple universities (SLR + case survey)	Mixed methods (SLR + case survey, CR/AVE tested)	Agile IT governance constructs (alpha 0.82-0.85) associated with faster decision-making and strategic alignment when supported by leadership, culture, and adaptive policy	5/5 (100%)
[97]	2025	Morocco; 110 staff	Quantitative survey (PLS-SEM)	Full text (N=110 Moroccan university staff, PLS-SEM) tests COBIT/ITIL-informed ITG mechanisms -> absorptive capacity -> digital transformation performance	3/5 (60%)
[128]	2025	USA; 4 universities	Qualitative multi-case study	IT project decisions correspond to ITG framework components (alignment, value, risk, performance)	4/5 (80%)
[17]	2025	Higher education; Web of Science 2023-2024	Scoping review	Four clusters: guidelines/frameworks, adoption challenges, stakeholder perceptions, research gaps	8/11 (73%)
[119]	2025	Public-sector financial system cloud migration	Case study & architectural evaluation	Evaluates hybrid service-oriented and domain-driven legacy decomposition on a government financial system migration (modularity 0.32 to 0.61) under data-integrity and compliance constraints	4/5 (80%)
[137]	2025	Iran; public universities (qualitative + N=63 leadership questionnaire)	Mixed methods (framework synthesis + N=63 survey)	Identifies resistance to change, fragmented technology, and leadership gaps; 63 university respondents rate feasibility of a human-centered digital transformation framework	4/5 (80%)
[140]	2025	India (Rajasthan); university administrative departments	Mixed methods (multi-department employee survey)	Multi-department university employee survey finds RPA improves administrative efficiency when paired with change management and cross-departmental standardization	4/5 (80%)
[85]	2025	South Africa; 3 provinces	Comparative case (policy + interviews)	Full text compares 3 South African provincial administrations using COBIT and ISO/IEC 38500 lenses; contrasts leadership-driven ICT governance progress against compliance-driven bureaucratic delay	5/5 (100%)
[148]	2025	EU; documentary analysis	Documentary analysis, qualitative coding	Principles become operational only as auditable controls with ownership and indicators across selection-to-retirement lifecycle	3/5 (60%)
[123]	2025	Kenya; public & private universities (cluster sample)	Quantitative/mixed survey & model design	Finds generic cloud security inadequate for Kenyan universities regulatory and infrastructure context; public universities prioritize scalability while private emphasize usability	4/5 (80%)
[73]	2025	Brazil; 3 federal universities (documentary)	Qualitative comparative case study	COBIT-adapted diagnostic scores ITG maturity 76.67% / 63.33% / 0% on public documents; volunteer practices absent; documents may not match practice	5/5 (100%)
[139]	2025	Oman; postsecondary HEIs (N=224 strategic-planning staff)	Quantitative survey (PLS-SEM)	Among 224 strategic-planning participants (83.9% response), both workflow automation (beta=0.36, p<0.001) and digitalization (beta=0.30, p=0.01) positively relate to strategic planning effectiveness	4/5 (80%)
[74]	2025	Global HE administration; 60 studies 2019–2025	Systematic literature review (PRISMA)	Administrative digital transformation is under-studied relative to teaching tech; RPA/AI efficiency depends on culture and strategic governance	9/11 (82%)
[188]	2023	China; 1 university	Case study (SWOT)	Examines administrative digitalization at one university (specific findings not extracted from abstract)	5/5 (100%)
[142]	2022	Public sector	Quantitative model test	Digital transformational leadership and agility positively associated with digital transformation	3/5 (60%)
[89]	2022	Brazil / international HEIs (literature + expert DSR evaluation)	Design science research (IT governance model)	DSR study synthesizes structure, process, and relational mechanisms (referencing COBIT and ISO/IEC 38500) into an evaluated IT governance model tailored to university administration	6/6 (100%)
[16]	2022	Public administration	Systematic review	Literature points from top-down hierarchical governance toward inclusive policy-making for actionable AI ethics	7/11 (64%)
[13]	2021	Global	Structured literature review	Nine barrier categories (22 barriers) incl. vision/strategy/policy, resources, leadership, skills, technology, resistance, government/economy	6/11 (54%)
[92]	2020	South Africa; case studies	Case study research (model proposal)	Full text (South African low-resource university case) references COBIT, ISO/IEC 38500, and ITIL; proposes context-fit ITG committee model for fragmented portfolios	4/5 (80%)
[189]	2019	Higher education; multiple cases	Multiple case study	Culture analyzed as a factor influencing IT-governance implementation	4/5 (80%)
[5]	2018	Not stated	Quantitative correlational	Intolerance of ambiguity related to resistance to change	3/5 (60%)
[88]	2017	6 case studies	Multiple case study	Full text (6 university cases) identifies baseline structure, process, and relational ITG mechanisms referencing COBIT, ITIL, and ISO/IEC 38500	4/5 (80%)
[87]	2016	Ghana; 1 public polytechnic	Case study (questionnaire, interviews)	Software procurement bypassed strict practice; inadequate funding; weak contracts/SLAs	4/5 (80%)
[129]	2024	Canada; 84 federal/provincial AI governance initiatives	Semi-systematic review / content analysis	Among 84 initiatives, programs and plans concentrate on industry/innovation and public administration; ethics standards, workforce services, and digital infrastructure receive less attention; outcome disclosure and public participation remain weak	9/11 (82%)
[106]	2026	Italy; Autonomous Province of Trento (30 staff; 6-month production)	Production HITL multi-agent deployment case	CARE multi-agent system over >160k administrative documents reports ∼70% drafting-time reduction on council inquiries while retaining mandatory human validation before publication	4/5 (80%)
[130]	2026	EU and Türkiye public administration (comparative)	Comparative institutional-legal analysis	EU AI Act risk-based institutional capacity contrasts with Türkiye's fragmented sectoral arrangements; gaps include missing formal risk classification, unclear accountability, and weak supervisory capacity	6/6 (100%)

TABLE XXIX
CONSTRUCTED TEST SCENARIO ASSUMPTIONS AND RA-GO V2 OUTPUT. S: VALUE SCALE (M$); r_f: RISK FACTOR; R: READINESS; θ_0: BASELINE COORDINATION; h: REGULATORY HEAVINESS (MINIMUM REVIEW INTENSITY $0.30/0.45/0.60$; $h{+}1$ MANDATORY SIGN-OFFS); REC.: MODAL RA-GO V2 ALTERNATIVE; $\bar{x}$: MEAN INTENSITY; TOP ACC.: MEAN TOP ACCEPTABILITY; M3: MODAL GAIM-DF ALTERNATIVE; BLOCKED: ALTERNATIVE EXCLUDED BY MANDATORY GATES

ID	Decision context	S	r_f	R	θ_0	h	Rec.	$\bar{x}$	Top acc.	M3	Blocked
A1	Enterprise generative-AI assistant for administrative and advising work	12	1.6	0.55	0.35	3	pilot/redesign	0.86	0.92	pilot	fast
A2	AI triage tool influencing admissions or financial-aid decisions (high-impact)	8	2.2	0.45	0.30	3	pilot/redesign	0.86	0.96	pilot	fast
A3	AI feature added to an already-approved SaaS product	3	1	0.65	0.50	2	adopt/redesign	1.00	0.94	adopt	fast
A4	Learning-analytics / AI tutoring vendor selection	6	1.4	0.50	0.40	3	pilot/redesign	0.80	0.99	pilot	fast
A5	Custom AI model built with a partner on institutional data	15	2	0.40	0.30	3	pilot/redesign	0.80	0.86	pilot	fast
B1	Student information system replacement	40	2	0.50	0.30	3	wrap/redesign	0.98	0.99	wrap	patch
B2	Finance / ERP modernization	30	1.7	0.55	0.35	2	wrap/redesign	0.87	0.98	wrap	patch
B3	HR / payroll: retain legacy vs replace	18	1.5	0.45	0.30	2	wrap/redesign	0.83	0.99	wrap	patch
B4	Research-administration system: interface wrapping vs rebuild	6	1	0.60	0.45	2	wrap/redesign	0.99	0.94	phased	patch
B5	Data warehouse and reporting platform consolidation	10	1.2	0.60	0.40	2	wrap/redesign	1.00	0.95	wrap	patch
C1	Third-party cloud vendor security assessment policy	5	2	0.60	0.45	3	tiered/redesign	1.00	0.91	tiered	attest
C2	FERPA school-official vendor data-sharing review	3	1.8	0.60	0.40	3	tiered/redesign	1.00	1.00	tiered	attest
C3	Accessibility conformance for purchased web and mobile software	8	1.3	0.50	0.35	3	tiered/redesign	0.90	0.85	remed	attest
C4	Inventory and contest path for AI-assisted advisory decisions	4	1.9	0.40	0.30	3	tiered/redesign	0.80	0.98	tiered	attest
C5	Safeguards oversight of service providers for student financial data	7	1.8	0.55	0.40	3	tiered/redesign	1.00	1.00	tiered	attest
D1	Redundant approval stages in a modernization acquisition workflow	2.5	0.8	0.60	0.30	2	redund/redesign	0.92	0.90	redund	drop
D2	Fragmented decision rights across IT, security, procurement, legal	4	1.1	0.45	0.25	2	central/redesign	0.70	0.75	redund	drop
D3	Handoff-heavy cross-functional approval path	3	0.7	0.65	0.35	1	redund/redesign	0.90	0.96	redund	drop
D4	Governance committee cadence and organizational inertia	2	1	0.50	0.30	2	redund/redesign	0.81	0.96	redund	drop
D5	AI-assisted routing and triage of review tickets (advisory only)	1.5	1.2	0.50	0.40	3	redund/redesign	0.80	0.62	redund	drop