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.

Governing work delegated to AI through dynamic control fit

Kevin McDermott

kmcdermott@wlu.ca

Wilfrid Laurier University Lazaridis School of Business and Economics https://orcid.org/0000-0002-4437-3779

Systematic Review

Keywords: AI delegation, agentic AI, organizational control, AI governance, systematic review, dynamic control fit

Posted Date: November 2nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11556449/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.

Governing work delegated to AI through dynamic control fit

Kevin McDermott

Lazaridis School of Business & Economics
Wilfrid Laurier University
Ontario, Canada

ORCID: https://orcid.org/0000-0002-4437-3779

Corresponding author

Kevin McDermott | kmcdermott@wlu.ca

Abstract

When organizations delegate work to artificial intelligence (AI), controls that remain formally in place may no longer achieve their purposes. We combine systematic identification and selection with integrative analysis of 171 papers across five research streams and 26 supplementary sources. The synthesis develops dynamic control fit: an arrangement's ability to achieve a specified control function while sustaining or restoring the conditions on which that function depends. We explain this ability through a framework of three interconnected requirements: Deliver, Renew and Adapt. Deliver concerns whether coverage, information, expertise and opportunities to intervene support effective control in current work. Renew concerns the practice, feedback and maintenance through which these capabilities persist. Adapt concerns recognizing deficiencies, obtaining authorized changes and verifying restored performance. The framework explains how AI may improve immediate enforcement while reducing experiences through which people develop judgment for future oversight, although other activities may compensate. Coverage, reviewer capacity and revision authority impose distinct constraints requiring different responses. Continued reliance must also be assessed separately from performance because beliefs, external expectations and practical constraints can sustain or limit use. By following a specified function across changes in control arrangements, the review develops testable explanations of improvement, erosion, compensation and recovery.

Keywords: AI delegation; agentic AI; organizational control; AI governance; systematic review; dynamic control fit

1 Control after delegation to AI

When organizations delegate work to artificial intelligence (AI), they change more than who performs particular tasks. They may also change how work is supervised, how errors are detected and who can intervene when something goes wrong. Consider an organization that uses AI to process payment requests. Its approval policy may remain unchanged, and software may reliably prevent payments that exceed an established threshold. Yet employees responsible for reviewing unusual transactions may have less information about how those transactions were processed, fewer opportunities to develop relevant expertise or insufficient time to intervene before a payment is completed. The organization has retained its formal controls, but its ability to achieve some of their intended purposes may have changed. The continued existence of a control does not establish its continued effectiveness.

We use the term *control function* to describe the purpose a particular control is intended to achieve, such as ensuring that a payment is authorized, detecting errors before a decision takes effect or resolving cases that require professional judgment. Organizations may delegate work to AI systems that recommend decisions, perform individual tasks or coordinate several actions. Agentic AI extends these possibilities by allowing systems to select and execute steps using delegated permissions, often through connections to other software tools. As organizations introduce these systems into their workflows, they can change who performs an activity, who evaluates its results and who has authority to interrupt or correct it (Baird and Maruping 2021). Such changes can improve existing controls, introduce new forms of oversight or weaken the arrangements on which effective control previously depended.

Organizational control theory provides an established foundation for examining these changes. Organizations direct and evaluate work through formal rules, supervisory review, performance measures, professional norms and other arrangements. These controls depend on particular conditions, including sufficient information about the work, knowledge of how it should be performed and the social relationships that support appropriate conduct (Ouchi 1979; Kirsch 1996). A supervisor cannot reliably evaluate work without some understanding of the task or a meaningful measure of its results. Similarly, professional norms depend on opportunities for people to develop and apply a shared understanding of acceptable practice. Delegating work to AI can alter these conditions even when an organization's formal policies remain unchanged. The effects vary: some controls may become more accurate or consistent, while others become less effective as information, expertise and authority are redistributed. Earlier research on automation illustrates the significance of these changes. In their examination of Australia's Robodebt scheme, Rinta-Kahila et al. (2022) document how the removal of manual accuracy checks altered verification responsibilities and contributed to failures in administrative decision-making. Although contemporary AI systems differ from the technologies examined in earlier automation research, such studies demonstrate why changes in the organization of work must be considered alongside changes in technical capability.

A growing body of research examines the governance challenges associated with organizational AI. Studies of AI governance identify the responsibilities, policies and organizational practices needed to direct and oversee these systems (Birkstedt et al. 2023; Hillebrand et al. 2025; Papagiannidis et al. 2025). Research on organizational control explains how organizations combine different controls and adjust their use as circumstances change (Cardinal et al. 2004, 2017). Delegation research examines how authority is assigned, how the capabilities of those performing work are assessed and how their activities are coordinated (Baird and Maruping 2021). Other scholars investigate the effects of automation on human expertise, the maintenance of AI systems, accountability for outcomes and the practical ability of people to intervene when problems arise (Bainbridge 1983; Nagbøl et al. 2024; Grote et al. 2026). Together, these perspectives provide important explanations of the organizational arrangements through which AI is governed.

These bodies of research examine different aspects of control and assess effectiveness using different evidence. Technical studies test whether a software safeguard detects or prevents specified actions under defined conditions. Human-oversight studies assess judgments, willingness to challenge AI recommendations and felt responsibility for decisions. Organizational governance studies examine responsibilities, policies and reporting

arrangements. Each perspective contributes relevant knowledge, but their outcomes differ. A software safeguard may operate accurately while failing to address some actions or omissions relevant to the organization's requirements. A person may retain formal responsibility for a decision while lacking the information or opportunity to exercise effective oversight. Similarly, the existence of a governance policy does not establish whether the activities performed under that policy achieve their intended results. Explaining effective organizational control therefore requires connecting these different forms of evidence to the particular functions that controls are expected to perform.

Recent reviews connect AI's technical arrangements with organizational roles and governance requirements. Schulte and Kanbach (2026) examine generative AI's roles in organizational decision-making, while Jeppe et al. (2026) examine tensions in algorithmic management. Helal et al. (2026) explain how perceived fidelity to an authorized AI delegation influences evaluation and reliance. Our outcome is different: the independently assessed performance of a specified control function as its supporting conditions change. Consider two AI-assisted payment workflows that resolve today's exceptional transactions equally well. In one, reviewers continue to practise difficult cases, receive feedback and can secure changes to deficient rules; in the other, relevant practice disappears and revisions depend on actors beyond their reach. Equal current performance—and even similar confidence in the systems—does not establish equal capacity to sustain or restore the authorization function. Existing work already explains learning, compensating activity and relocated decision authority (Ertem 2026; Shoaib et al. 2026; Kumar and Suthar 2026). Our contribution is to connect these mechanisms through a common functional benchmark, distinguishing present effectiveness, renewal and authorized correction from continued reliance.

This problem requires attention to both the immediate effectiveness of a control and the conditions that allow it to remain effective. A software check may improve the detection of routine errors while reducing the opportunities through which employees develop the expertise needed to recognize more complex problems. An organization may have knowledgeable reviewers today but lack arrangements for sustaining that expertise as employees leave or tasks change. Alternatively, staff may recognize that an AI system no longer applies an appropriate rule but lack the authority to modify it or obtain a revision from its provider. These situations involve different weaknesses in organizational control. They also demonstrate why assessing a control at one point in time is insufficient to explain whether it can continue to achieve its purpose. An adequate explanation must connect present performance with the activities that sustain its supporting conditions and the organizational capacity to respond when those conditions become inadequate.

We therefore address one overarching research question: *What conditions allow organizational control functions to remain effective as work is delegated to AI, and how can those conditions be sustained or restored over time?*

To address this question, we combine systematic identification and selection with integrative analysis across research traditions. The main corpus comprises 171 papers across five streams: organizational AI governance, delegation and decision rights, runtime enforcement, human roles and accountability, and managerial control delegated to AI. We use 26 supplementary sources to examine alternative explanations and refine the interpretation. Comparing functions, their supporting conditions and the evidence used to assess them allows us to connect current performance with the activities and authority needed to sustain or restore it. The streams organize this comparison; they do not presume that different measures demonstrate equivalent outcomes.

The central theoretical contribution is the concept of *dynamic control fit*, which we define as an arrangement's ability to meet the requirements of a specified control function over time while sustaining or restoring the conditions on which that function depends. This concept shifts attention from the continued presence of particular control arrangements to their capacity to achieve an intended purpose as the organization of work changes. Replacing a control’s mechanism, target or performer does not by itself eliminate its function. Conversely, an unchanged arrangement may become less effective because the information, expertise or authority on which it depends has deteriorated. Following a specified function across these changes allows researchers to distinguish the reconfiguration of control from changes in its actual performance.

Our synthesis identifies three interconnected requirements for achieving and sustaining dynamic control fit. **Deliver** concerns whether a control can achieve its intended function under current conditions. This depends on whether the arrangement addresses the relevant actions and omissions and whether people or systems have sufficient information, expertise and opportunities to intervene. **Renew** concerns the activities that preserve these conditions over time, including continued practice, feedback, maintenance and complementary work that sustains organizational capabilities. **Adapt** concerns the ability to recognize when a control no longer meets its requirements, obtain authorized changes and establish whether those changes restore effective performance. These requirements are closely related rather than sequential stages. For example, an organization may need to maintain employee expertise while simultaneously revising a software safeguard that no longer addresses the work being performed.

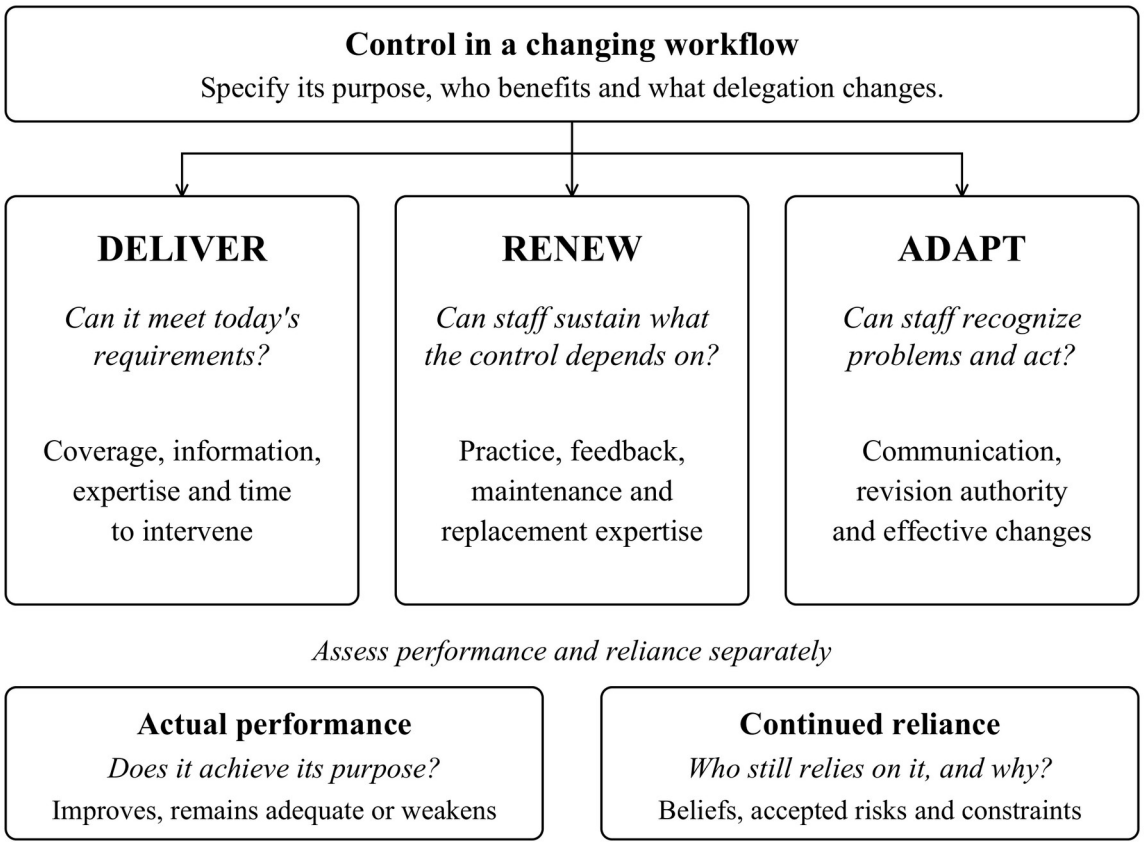
Together, these requirements explain how improving one aspect of control can create vulnerabilities elsewhere. Research on automation and expertise shows how the redistribution of work can affect opportunities to develop judgment, exercise discretion and maintain the ability to evaluate later decisions (Ertem 2026; Kumar and Suthar 2026; Shoaib et al. 2026). Delegating routine decisions may improve consistency and reduce immediate demands on reviewers while changing the experience available to those who must eventually resolve unusual cases. Whether this produces a decline in control depends partly on the presence of other activities that sustain expertise and provide opportunities for effective review. Similarly, the ability to detect a problem does not guarantee the ability to correct it when decision rights or technical authority have shifted to other actors. Dynamic control fit connects these immediate and longer-term relationships within a common theoretical explanation.

The synthesis also distinguishes three constraints on effective control: *coverage*, *reviewer capacity* and *revision authority*. Coverage concerns whether an arrangement addresses the actions and omissions relevant to its intended function. Reviewer capacity concerns whether people have the knowledge, information, attention and time needed to evaluate cases referred to them. Revision authority concerns whether those who identify a problem can implement an appropriate change or obtain action from someone authorized to do so. These constraints help explain why different weaknesses require different organizational responses. Extending a software check does not by itself replace the human expertise needed to judge exceptions; additional training cannot overcome the absence of authority to correct a deficient rule.

Finally, our explanation distinguishes the actual performance of a control from an organization's continued reliance on it. Organizations may retain arrangements because they believe those arrangements are effective, accept known limitations or lack workable alternatives. Continued use therefore cannot, by itself, establish that a control continues to achieve its intended purpose. This distinction is particularly important when AI changes the visibility of work or separates responsibility for outcomes from the capacity to assess and correct them. The framework supports investigation of not only how controls become more or less effective, but also why organizations recognize, respond to or continue operating with those changes.

Figure 1 introduces dynamic control fit and Deliver, Renew and Adapt. Section 2 explains the organizational control and delegation foundations; Section 3 describes identification, selection and integrative analysis. Section 4 reports the corpus mapping and develops the Deliver, Renew and Adapt explanation and then considers functional trajectories and continued reliance. Section 5 specifies empirical tests, Section 6 considers organizational assessments, and Section 7 concludes.

Fig. 1 Sustaining effective control after AI delegation



Deliver, Renew and Adapt express related requirements for achieving and sustaining dynamic control fit for a specified function.

2 Delegation and organizational control

2.1 Assigning work and retaining practical authority

Delegating a task to AI determines what the system may perform, but not whether the organization retains the means to direct, evaluate and correct that work. In a payment workflow, for example, managers may permit AI to prepare and process requests while reserving approval of exceptional transactions for employees. Effective governance then depends on which actions require authorization, what information reaches the employees and whether intervention is possible before a transaction takes effect. Baird and Maruping (2021) explain delegation through capability assessment, assigned decision rights and continuing coordination. These remain distinct considerations after a system begins operating, especially as its assigned tasks and opportunities for intervention change.

Formal responsibility is not evidence of practical control. An employee can retain approval authority yet routinely accept an AI recommendation, lack sufficient information to evaluate it or use an override for an objective other than improving that decision (Kim et al. 2024). Conversely, software permissions can prevent an AI system from exercising an action it is technically capable of executing. Accountability requires a feasible means of evaluating or influencing outcomes, rather than simply naming someone responsible (Grote et al. 2026). These distinctions apply both when AI performs operational tasks that people supervise and when AI itself allocates or evaluates employees' work.

These decisions can vary within a single workflow. The organization might allow an AI system to reject requests that violate a fixed rule, while requiring a person to approve cases with conflicting information. It may also make the person accountable for reporting repeated system errors without providing access to modify the

rule. The organization thus retains several kinds of authority—execution, review, and revision—that need not rest with the same actor. Treating them as a single "human in the loop" arrangement conceals the differences relevant to control performance.

2.2 Conditions and dependencies of organizational control

Organizational control theory provides an account of what makes control possible. Behaviour controls require sufficient knowledge of how work should be performed, while outcome controls require meaningful measures of its results (Ouchi 1979). Other controls rely on professional expectations, socialization and individuals' ability to regulate their own conduct (Kirsch 1996). Each form therefore depends on information, evaluative knowledge, opportunities to intervene or social relationships. Detailed activity logs, for example, may make an AI system's actions visible without giving a reviewer the expertise needed to judge them. A change in technology alters particular prerequisites, not necessarily the effectiveness of all controls in the workflow.

Controls also operate in combinations that organizations adjust over time (Cardinal et al. 2004, 2017; Schafheitle et al. 2020). Reviewing exceptions can both prevent current mistakes and teach employees how to recognize difficult cases later. An additional automated referral rule may help one function while overloading the people responsible for another. Such relationships mean that an apparent improvement to an individual check may depend on, or weaken, activities outside that check. Bainbridge (1983) identifies the longstanding problem of reserving difficult interventions for employees whose opportunities to practise have diminished. Organizational testing, maintenance and alternative assignments can preserve capabilities that the immediate arrangement no longer develops (Asatiani et al. 2021; Nagbøl et al. 2024; Ertem 2026).

2.3 Assessing functions when arrangements change

A control function is the specific purpose to be achieved, such as authorizing a payment or resolving an exceptional case. It may be supplied by a person, software, or a combination of both. Reconfiguration changes the mechanism, performer, target or timing of a control; it does not by itself establish that the function has improved or deteriorated. Substituting a software check for human approval may strengthen routine authorization while transferring the authority to revise its rules to an external provider. Functional assessment must consequently examine what the arrangement accomplishes and whether the prerequisites of that performance remain available.

Functional erosion occurs when performance of the specified function declines; functional loss occurs when it ceases without an effective replacement during the period examined. Neither should be inferred simply from the removal of a familiar procedure. Other activities may compensate for a lost prerequisite, and an unchanged procedure may fail to achieve its purpose. Bromley and Powell (2012) distinguish activities formally performed from the ends they attain.

The time period matters when deciding whether performance has changed. If a safeguard reliably prevents prohibited actions this month but the organization is withdrawing the maintenance work necessary to keep its rules current, the evidence supports current effectiveness and a possible future vulnerability, not yet observed erosion. If a weakness later emerges and is corrected, the interval during which the control failed should remain visible rather than being treated as uninterrupted success. Distinguishing these possibilities prepares the temporal analysis without assuming the answer in advance. Continued reliance cannot settle the matter either: decision makers may misunderstand deteriorating performance, knowingly accept a limitation or lack an alternative. Nor does success against a managerial objective settle whether the objective or its effects on workers are justified (Kellogg et al. 2020). Table 1 identifies how these established theoretical foundations inform the more specific explanation developed below.

Table 1. Theoretical foundations and their contributions to dynamic control fit

Theoretical foundation and sources	Established explanation	Contribution to dynamic control fit
Control prerequisites Ouchi (1979); Kirsch (1996)	Behaviour controls require knowledge of the work; outcome controls require meaningful measures of results. Social and self-controls draw on shared expectations and self-regulation. Visibility of an action is not equivalent to the knowledge needed to evaluate it.	**Deliver.** Identifies the information, expertise, evaluative criteria and opportunities for intervention required to achieve a specified control function. More extensive monitoring need not improve judgment if the necessary knowledge is absent.
Control configurations and dependencies Cardinal et al. (2004, 2017); Schafheitle et al. (2020)	Organizations combine and adjust formal, technical and informal controls. A change in one arrangement can support or interfere with another by altering information, workload, learning opportunities or other conditions on which it depends.	**Deliver / Renew.** Explains why improving one control can weaken or strengthen the current or future effectiveness of another. Compensating activities may preserve a function despite changes to the original arrangement.
Delegation and accountability Baird and Maruping (2021); Grote et al. (2026)	Delegation involves assessing an agentic artifact, assigning authority and coordinating its work. Responsibility and formal decision rights do not guarantee that responsible actors can evaluate actions, intervene or secure a correction.	**Deliver / Adapt.** Distinguishes authority on paper from the practical capabilities needed for oversight and from an effective route to authorized changes when a control becomes inadequate.
Expertise, maintenance and compensation Bainbridge (1983); Asatiani et al. (2021); Nagbøl et al. (2024); Ertem (2026)	Automation may reduce the practice through which people acquire expertise. Ongoing evaluation, feedback, technical maintenance and alternative assignments may sustain capabilities that are no longer reproduced through routine work.	**Renew.** Connects present performance with the continuing production of its prerequisites. Declining experience need not cause erosion if another activity supplies the required capability and sustains actual performance.
Decision rules and professional judgment Kumar and Suthar (2026); Shoaib et al. (2026)	Delegation can shift discretion from individual cases toward the rules, data and configurations governing repeated decisions. Changes in opportunities to reason through work can also affect later judgment and error detection.	**Renew / Adapt.** Explains why maintaining reviewer expertise and obtaining authority to revise decision rules are distinct requirements, and why corrective capacity can change when discretion moves.
Decoupling, reliance and contestation Bromley and Powell (2012); Kellogg et al. (2020)	Organizations may perform prescribed activities without achieving their intended outcomes. People subject to control can contest its aims, implementation or consequences. Continuing use of an arrangement is not proof of effectiveness.	**Assessment and boundary conditions.** Requires separate evidence about actual functional performance and continued reliance, and identification of the purpose, beneficiaries and consequences of the control being assessed.

Note. Connections to Deliver, Renew and Adapt identify how established explanations inform the synthesis; they do not imply that each cited study empirically tests dynamic control fit.

3 Search, selection and integrative analysis

3.1 Identification, eligibility and source groups

We combined systematic procedures for identifying and selecting literature with integrative analysis because research on organizational control after AI delegation spans several disciplines, each using different concepts, methods and forms of evidence (Torraco 2016; Elsbach and van Knippenberg 2020; Block et al. 2025). Rather than estimate the average effect of a particular intervention, our purpose was to develop an explanation of how different organizational and technical conditions contribute to effective control. Between 6 and 8 October 2026, we searched Web of Science Core Collection and ScienceDirect using combinations of terms relating to AI, delegation, organizational governance and control. The database searches covered articles and reviews without

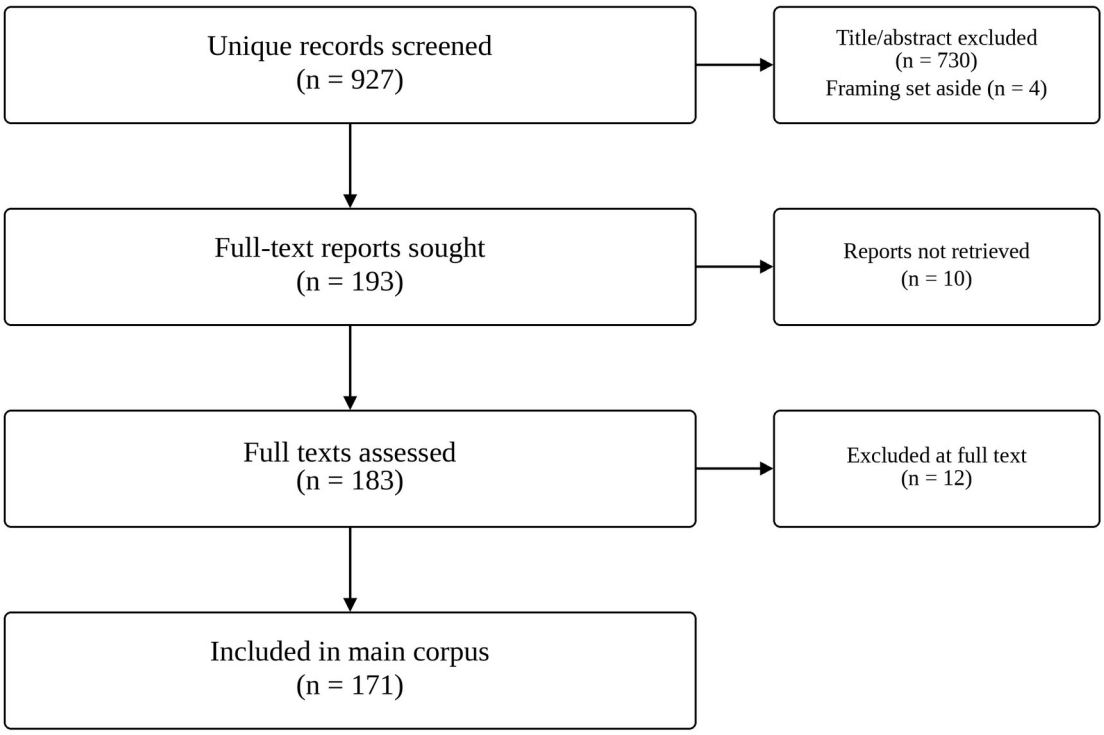
restrictions on publication year or language. Scholarly preprints, conference proceedings and book chapters identified through other routes were also eligible.

We complemented these searches by examining publications citing Baird and Maruping (2021) and identifying additional relevant literature manually. Further searches used OpenAlex to identify publications through citation relationships and related topics. We also examined one generation of references cited by Kellogg et al. (2020), Bainbridge (1983) and Nagbøl et al. (2024). Supplementary sources were selected to examine alternative explanations, challenge emerging interpretations and clarify theoretical relationships, rather than to provide a representative sample of the literature. Online Resource 1 describes the search coverage, eligibility boundaries and analytical procedures. Online Resource 2 documents the search procedures, discovery routes and selection decisions.

We removed duplicate publications using digital object identifiers (DOIs) and standardized bibliographic information. We then screened the titles and abstracts of 927 publications. Of these, 730 were excluded using criteria listed below, four were retained separately for theoretical framing, and 193 were sought for full-text assessment. To qualify, a publication had to examine AI-related work in an organizational setting and substantively address how that work was governed, authorized, evaluated, corrected or made subject to accountability. We excluded studies concerned only with AI adoption or technical performance when they did not establish a meaningful connection to organizational control. Technical studies were eligible when they examined how systems support or enforce organizational requirements through mechanisms such as permissions, referrals, validation or auditing. Research on earlier forms of automation could also inform the analysis, although we did not assume that its findings applied directly to contemporary AI agents capable of executing sequences of interconnected actions.

We closed full-text retrieval on 23 October 2026. Of the 193 publications sought, we obtained and assessed the full texts of 183. We could not obtain the remaining 10 through the university library, publisher websites or open repositories, so we did not assess them or count them as exclusions. Full-text assessment excluded 12 of the 183 publications because they did not meet the eligibility criteria, which left 171 publications in the main corpus. The 22 candidates that did not enter the corpus therefore comprise 12 substantive exclusions and 10 unavailable reports. Online Resource 1 reports this flow of publications, and the full-text decisions in Online Resource 2 record the outcome for each of the 193.

Fig. 2 PRISMA-style flow of publication selection for the main corpus



Note. The 26 supplementary synthesis papers are separate from the 171-paper main corpus.

3.2 Extraction, comparison and evidentiary limits

For each publication, we recorded the control function, earlier control arrangement, proposed or observed changes, relevant conditions, theoretical perspective, level and setting, evidence category, and type of AI. The 171 main-corpus publications each received one primary research-stream assignment; secondary stream associations did not affect stream totals. We examined the source sections relevant to the review question, retained locators for consequential claims and read particularly informative studies more closely. Reading depth differed by relevance and evidence rather than following an identical cover-to-cover protocol for every publication.

We compared publications by the function being controlled and the prerequisites on which its performance depended. Evidence of an activity—for example, an automated referral—was separated from the resulting outcome, such as correct and timely human resolution. We distinguished empirical observations from simulations, technical designs, conceptual arguments and reviews, considered the level of analysis, and examined whether later capability could depend on practice, feedback, maintenance, recognition and authorization. Comparisons also examined changes in task requirements, staffing and incentives as alternative explanations.

The synthesis compared supporting, conflicting and qualifying findings. It treats delayed expertise erosion, compensating maintenance and recovery after correction as conditional theoretical relationships unless a source directly observes the relevant sequence. This distinction separates the framework's explanatory propositions from measured effects in the publications reviewed.

We assessed evidentiary strength through study design, outcome measures, setting and observation period, including whether a reported result was independently observed or self-reported. Studies drawing on the same case, including Asatiani et al. (2021) and Nagbøl et al. (2024), were not treated as independent replications. Source identities, DOIs and versions were checked; consequential evidence appraisals are illustrated in Online Resource 2.

The author conducted eligibility assessment, extraction and interpretation. A sensitivity check excluding 19 less-directly relevant papers left 152 publications and is reported in Online Resource 1. The original searches, selective supplementary discovery, uneven reading depth and ten unavailable full texts limit completeness and generalizability. The evidence offers a basis for theory development, not a single longitudinal test of the complete framework.

Generative-AI tools assisted with selected drafting, preparation of tables and exploratory analysis of review records. The author assessed these outputs against the source literature and the recorded selections and takes responsibility for the final analysis and text.

4 Descriptive results and theoretical synthesis

4.1 Descriptive results: where the main corpus contributes

The main corpus contains 171 publications: 64 empirical studies, three simulations, 65 conceptual papers, 23 reviews and 16 design artifacts. The five primary research streams contain 34 organizational AI governance, 52 delegation and decision-rights, 23 runtime enforcement, 50 human roles and accountability, and 12 AI-directed managerial-control publications (Table 2). Twenty-six supplementary synthesis publications informed the interpretation but did not enter these descriptive counts.

To show how the corpus informed the proposed framework, we assigned each main-corpus publication one primary analytical emphasis based on the recorded control arrangement, supporting conditions and treatment in the synthesis. The categories distinguish two substantive concerns within each of Deliver, Renew and Adapt, alongside cross-cutting papers and broader contextual work. This is a transparent index of where publications contribute to the argument, not a claim that each paper proves its component or that a study addresses only one component. Online Resource 2 identifies the assigned category for every main-corpus paper in the Source inventory.

Table 2. Research streams, evidence and contributions to dynamic control fit

Primary stream and illustrative sources	Principal contribution to the review	Limits of evidence and relevance to dynamic control fit
Organizational AI governance *34 papers; 14 empirical or simulation* Mäntymäki et al. (2022); Birkstedt et al. (2023); Mökander et al. (2022); Nagbøl et al. (2024)	Explains how organizations assign oversight responsibilities, connect AI to established governance and conduct ongoing evaluation, validation and maintenance. Governance activities can contribute to current control and preservation of supporting capabilities.	A policy, accountable office or completed review does not demonstrate that the intended function was achieved. Deliver requires evidence about performance; Renew requires evidence that ongoing practices maintain the capabilities on which it depends.
Delegation and decision rights *52 papers; 11 empirical or simulation* Baird and Maruping (2021); Kim et al. (2024); Helal et al. (2026); Kumar and Suthar (2026)	Explains how decision authority, oversight and discretion are distributed between people and AI. Distinguishes formal responsibility from actual decision behaviour and examines how discretion can shift to software rules and configurations.	Formal decision rights do not establish independent judgment, effective intervention or the authority to revise a deficient arrangement. These distinctions connect Deliver to current oversight and Adapt to access to corrective authority.
Runtime enforcement *23 papers; 5 empirical or simulation* Zwerdling et al. (2025); Pacella et al. (2026); Cuneo et al. (2026); Ling et al. (2026)	Describes and evaluates restrictions, executable rules, permissions and referrals applied while AI systems act. Shows that different safeguards address different actions, omissions or stages of an organizational workflow.	Successful execution of specified checks does not establish complete coverage, effective downstream human review or sustained organizational outcomes. Deliver depends on coverage and subsequent intervention; Adapt may require revising technical rules.

Primary stream and illustrative sources	Principal contribution to the review	Limits of evidence and relevance to dynamic control fit
Human roles and accountability *50 papers; 33 empirical or simulation* Liel and Zalmanson (2025); Gräfe et al. (2024); Rinta-Kahila et al. (2022); Ertem (2026)	Examines how employees interpret, verify and correct AI-supported work, and the expertise, attention and often overlooked labour involved. Identifies how delegation can change opportunities to develop the judgment needed for later review.	Assigning a reviewer or demonstrating current task performance does not establish independent assessment or the future availability of expertise. Deliver depends on actual reviewer capacity; Renew depends on practice, feedback and compensating activities.
Managerial control delegated to AI *12 papers; 4 empirical or simulation* Kellogg et al. (2020); Hillebrand et al. (2025); Schafheitle et al. (2020); Leavitt et al. (2025)	Examines AI systems that allocate tasks, monitor or evaluate people. Shows why the performer, target and purpose of control must be distinguished, and why those affected may contest organizational objectives or procedures.	More consistent monitoring does not establish that a control objective is justified or its consequences acceptable. Assessing functional performance requires specifying whose purpose is served and the relevant obligations and interests.
Main coded corpus	171 papers: 64 empirical; 3 simulations; 65 conceptual; 23 reviews; 16 design artifacts.	Each main-corpus paper has one primary stream. Counts exclude supplementary and framing sources. Illustrations include Kumar and Suthar, Ling et al. and Ertem from the 26 supplementary synthesis sources, and Kellogg et al. for theoretical framing; other examples here are main-corpus papers. Online Resource 2 gives complete memberships.

Table 3. Primary analytical emphasis of papers in the main corpus

Requirement and primary analytical emphasis	Papers, n (empirical / simulation)	Illustrative main-corpus sources
Deliver · D1 Coverage / enforceable boundaries	28 (5)	Pacella et al. (2026); Cuneo et al. (2026); Zwerdling et al. (2025)
Deliver · D2 Human judgment / practical intervention	44 (20)	Liel and Zalmanson (2025); Kim et al. (2024); Rinta-Kahila et al. (2022)
Renew · R1 Learning / future expertise	8 (4)	Frieri and Orelli (2026); Lovett (2026); Naidoo (2026)
Renew · R2 Maintenance / compensating work	26 (17)	Nagbøl et al. (2024); Mökander et al. (2022); Kathuria et al. (2026)
Adapt · A1 Recognition of emerging mismatch	8 (3)	Tarafdar et al. (2026); van Esch and Cui (2026); Campbell (2026)
Adapt · A2 Corrective authority / revision route	14 (4)	Baird and Maruping (2021); Kuziemski and Misuraca (2020); Safaei et al. (2026)
Cross-cutting · reliance / legitimacy / impacts	31 (14)	Tiwari and Jain (2026); Hillebrand et al. (2025); Schafheitle et al. (2020)
Broad governance / delegation context	12 (0)	Birkstedt et al. (2023); Mäntymäki et al. (2022)
TOTAL · mutually exclusive primary emphasis	**171 (67)**	**Complete mappings: Online Resource 2, Source inventory, final column**

Note: One primary analytical emphasis per main-corpus publication. Parentheses show empirical studies or simulations within each category. Online Resource 2 gives all 171 source-level assignments; supplementary and framing publications are excluded.

Of the 171 papers, 72 have Deliver as their primary emphasis (28 on executable coverage and 44 on human judgment or intervention); 34 concern Renew (eight on learning and future expertise and 26 on maintenance or

compensating work); and 22 concern Adapt (eight on recognizing emerging mismatches and 14 on corrective decision rights and revision routes). Thirty-one primarily address reliance, legitimacy or other cross-cutting questions, while 12 provide broader context. These primary-emphasis counts reveal a marked imbalance between research on present control arrangements and research examining how their supporting conditions change over time.

The sparsest components concern precisely the temporal mechanisms that motivate this review. Learning and future expertise (R1) contains eight publications, only four empirical; recognizing emerging mismatch (A1) also contains eight, only three empirical. P3 asks whether reduced practice and feedback cause subsequent deterioration of reviewer judgment, while P5 concerns the consequences of an inadequate route from recognized deficiency to authorized correction; P5 also draws on the distinct A2 literature about revision authority. The central empirical gap identified by the review is the limited evidence following these processes over time—from lost opportunities to practice to independently assessed later expertise, or from emerging warning signs through recognition, authorization and verified correction. We therefore present P3 and P5 as testable conditional propositions.

4.2 Control functions and conditions for sustained effectiveness

The synthesis takes the specified control function, rather than a particular rule, technology or job role, as its unit of analysis. A function describes what an arrangement must accomplish, for whom, and with respect to which actions or omissions. Human review, automated rules or combined activities can supply the same function, while a workflow can contain several functions with different performance. In a hypothetical payment workflow, for example, authorizing ordinary transactions and resolving exceptional cases are separate control functions. A safeguard maintained by an external provider might apply an approval threshold, while an employee examines unusual requests. The organization must establish whether each function is achieved and which other activities supply the information, judgment and authority it requires.

Dynamic control fit describes the capacity to achieve a specified function over time while sustaining or restoring the conditions on which it depends. It distinguishes three related aspects that cannot be inferred from one another. Present effectiveness requires evidence that the function is achieved under existing requirements. Renewal capacity concerns the ongoing activities that preserve the conditions for subsequent performance. Adaptation capacity concerns the possibility of recognizing inadequacy and obtaining an effective correction. A control may perform well now while its renewal capacity deteriorates; a failing control may remain recoverable. Recovery after a failure, however, is not uninterrupted effective control. These are distinct empirical claims, even though a single activity can contribute to more than one.

A review meeting illustrates the distinction. It may resolve an exceptional case, give participants feedback useful for later judgments and expose an approval rule that needs revision. These are three different contributions to control: present execution, renewal of capability and identification of a need to adapt. Following the same specified function through subsequent decisions connects them across time.

Organizational control theory explains why information, measurement and authority are prerequisites of effective control (Ouchi 1979; Kirsch 1996; Cardinal et al. 2004, 2017). The synthesis extends this reasoning to activities that replenish these prerequisites and to the routes through which organizations can correct deficient arrangements. The next sections explain the relationships shown by the corpus evidence.

4.3 Deliver: Coverage, judgment and opportunities for intervention

Deliver concerns whether an arrangement can achieve its intended function in current work. The first requirement is coverage: a control must address the relevant actions and omissions, not merely execute its programmed checks reliably. The payment safeguard might reject a proposed transaction above a permitted threshold yet fail to detect that an approval request was never made. The software could operate correctly on every action presented to it while the organization's authorization function remained incomplete. Control

adequacy therefore cannot be inferred from technical reliability without specifying the wider requirement and the actions that the mechanism is capable of observing.

Runtime-enforcement architectures use permissions, restricted tool access and executable policy rules to constrain AI actions (Cuneo et al. 2026; Jørgensen and Ma 2026; Arora et al. 2026; Zheng et al. 2026). Zwerdling et al. (2025, pp. 600–601) report a pass¹ benchmark score rising from 0.450 to 0.685 across 22 policy-violating tasks in a simulated airline workflow after adding enforcement. The score assesses agreement between final and intended database states, rather than measuring policy compliance alone. It supports improved task performance under the tested arrangement. The architecture inspects proposed tool calls; an omitted required action need not generate a call for inspection. Comprehensive organizational coverage therefore does not follow from that result. Nor should this observability limit be generalized to every possible runtime safeguard: tested outcomes and the boundary of a particular check remain separate questions.

The work following a referral introduces a second requirement. Pacella et al. (2026) test 166 constructed profiles against explicit escalation criteria. A fixed rule corrects missed referrals by Qwen2.5, while another evaluated model refers relevant cases but also sends excessive additional cases. Missed referrals create a coverage problem because cases that require human attention never reach it. Excessive referrals raise a different possibility: a referral mechanism may work according to its rules but place demands on employees that exceed their ability to resolve exceptions. Referral and resolution are distinct functions. The study tests the former under its specified conditions; it does not demonstrate the latter in an operating organization. Our theoretical inference is that a more active or accurate automated control can undermine a related human control if it changes downstream workload without sufficient review capacity.

Governance studies reveal why organizational responsibility alone does not close these gaps. AI governance research describes policies, accountable offices and the incorporation of AI into established oversight processes (Mäntymäki et al. 2022; Mökander et al. 2022; Birkstedt et al. 2023; Papagiannidis et al. 2025). Kathuria et al. (2026), for example, report how dispersed software-bot development increases demands on central approval and monitoring teams. Such arrangements can specify who should examine an exception. They do not establish whether that person sees the relevant case, understands it, or can act before the outcome occurs. Delegation research similarly separates the assignment of authority from its exercise: Baird and Maruping (2021) discuss delegated rights and coordination, and Yan et al. (2026) consider supervision in relation to reversibility and verifiability. A retained human approval right is a design feature, not direct evidence of the function it achieves.

The distinction is particularly important in agentic workflows, where an AI system may initiate several linked actions before the record reaches a supervisor. A check placed at the first request, a permission limiting tool execution and a later review of completed transactions have different opportunities to prevent harm. None is categorically superior: the appropriate arrangement depends on the function, the reversibility of the action and whether its consequences can be evaluated. What an organization labels "oversight" must be translated into the activity and result being assessed.

The ability to evaluate a case depends on information that is meaningful to a person with relevant expertise. Organizational control theory distinguishes observing behaviour from understanding whether it is appropriate, and measuring outcomes from knowing whether a measure represents the desired result (Ouchi 1979; Kirsch 1996). A detailed log can show which step an agent executed without revealing whether that step was substantively justified. Lebovitz et al. (2021) illustrate how labels used to assess AI outputs can fail to represent contextual professional judgments. Ante (2026) and Basu (2026) identify difficulties in specifying objectives or values that have previously been interpreted implicitly. These accounts suggest that adding monitoring data cannot by itself supply the interpretive knowledge required for adequate evaluation.

Human-AI judgment research cautions against equating a reviewer with effective review. In four image-classification experiments with 1,445 crowdworkers, Liel and Zalmanson (2025) find that erroneous AI advice can influence judgments despite sufficient task skill. Kuziemski and Misuraca (2020) describe Poland's unemployment profiling system, where clerks retained discretion but, according to a cited secondary study, questioned fewer than one in 100 categorizations. That rate was not measured by these authors and cannot itself

establish decision accuracy. Together, the accounts distinguish formal authority from its exercise under social and organizational constraints. Kim et al. (2024) likewise show that overrides may serve objectives other than decision accuracy. Independent judgment and successful error detection therefore require assessment beyond the presence of a review step.

Timing and reversibility complete the account of current control. A reviewer may identify a mistake after an irreversible payment and support later accountability, but cannot prevent that original outcome. Pre-execution authorization, prompt exception resolution and retrospective auditing are different functions even when they rely on some of the same records. For high-impact or irreversible actions, the available interval for intervention may be decisive. Deliver therefore depends on appropriate coverage, evaluative information and expertise, and an opportunity for intervention suited to the function. These prerequisites are not interchangeable. More logging does not correct a missing approval requirement; training alone cannot prevent an action completed before review; and improved referral does not ensure that an overloaded reviewer can resolve cases.

A change in delegation strengthens current control when it supplies an unmet prerequisite and the other prerequisites remain adequate. It can instead weaken a related function when it transfers work to an already constrained reviewer. Table 4 compares the available evidence and the resulting cross-stream interpretations.

Table 4. Cross-stream evidence informing the requirements for dynamic control fit

Analytical concern	Representative evidence and its limits	Cross-stream inference for dynamic control fit
Coverage of the required function (Deliver)	Zwerdling et al. (2025) improve benchmark task success under policy enforcement across 22 simulated airline tasks; final-state agreement is not a standalone compliance measure. Checks of proposed calls need not capture omitted required actions. Governance studies describe oversight practices (Mökander et al. 2022; Nagbøl et al. 2024), not equivalent coverage tests.	Establishing a governance arrangement or demonstrating that a safeguard enforces specified rules does not show that all actions and omissions relevant to the control function are addressed. Coverage must be evaluated against the function's requirements and the actions the arrangement can actually observe.
Referral versus resolution (Deliver)	Pacella et al. (2026) test escalation decisions for 166 constructed profiles, distinguishing missed referrals from unnecessary ones. Liel and Zalmanson (2025) find that erroneous AI advice can influence human judgments in four image-classification experiments. These studies evaluate different tasks and do not track referrals through one organizational workflow.	A safeguard that identifies cases for review has not thereby shown that people can resolve them. How technical referral affects control depends on the composition and volume of cases, reviewers' information, expertise and available time. Joint organizational effects remain to be examined.
Responsibility and practical control (Deliver / Adapt)	Kim et al. (2024) show that overrides can serve objectives other than decision accuracy. AlQershi et al. (2026) describe accountable leadership under constrained discretion using a survey and interviews. Their self-reported evidence does not independently establish actual override behaviour or restored control.	Formal responsibility, observed override behaviour and the capacity to assess or change decisions are distinct. Effective oversight requires evidence of what an accountable actor can actually evaluate, interrupt or correct; assigning decision rights alone cannot establish those capabilities.
Future supply of judgment (Renew)	Bainbridge (1983) and Ertem (2026) explain how automation can change learning opportunities. Frieri and Orelli (2026) distinguish productivity from prospective learning effects. Naidoo (2026) links initial reliance to later reported complacency among industrial customers, without independently assessing reviewer expertise.	A control may work now because experienced people remain available even while the activities that develop future reviewers are changing. Whether capacity erodes depends on the work retained, feedback and replacement opportunities; immediate productivity gains do not establish a later loss of expertise.

Analytical concern	Representative evidence and its limits	Cross-stream inference for dynamic control fit
Maintenance and compensation (Renew)	Nagbøl et al. (2024) describe quality discussions, human validation and maintenance, while Asatiani et al. (2021) examine related practices surrounding AI use. The studies have overlapping authors and concern related work in the same organizational setting; they are connected rather than independent empirical confirmations.	Effective control may depend on maintenance and complementary activities outside the software safeguard itself. Such activities can preserve a needed capability when earlier work disappears, but their presence is not sufficient: subsequent capability and functional performance must be established.
Recognition and corrective capacity (Adapt)	Kumar and Suthar (2026) locate discretion in configurations governing repeated decisions. Kuziemski and Misuraca (2020) trace complaints, external scrutiny and withdrawal of Poland's profiling tool. This demonstrates a correction route, without measuring the subsequent recovery of the service function.	Recognizing or reviewing a problem is not equivalent to obtaining a correction. Recovery requires a workable route to appropriately authorized intervention and evidence that the change restores the specified function, including attention to consequences for other controls.
Actual performance versus continued reliance (Cross-cutting)	Tiwari and Jain (2026) find that advertising practitioners' interpretations of client expectations can limit acted-upon reliance despite performance trust; client positions were not directly verified. Bromley and Powell (2012) separate prescribed means from achieved ends. Neither use nor trust alone measures functional performance.	Assess functional performance separately from reliance decisions. Investigate beliefs, attributed and verified external expectations, accepted limitations and alternatives. Illusory control persistence requires mistaken confidence; externally restricted use or knowingly retaining an inadequate arrangement has a different explanation.
Purpose and affected participants (Assessment boundary)	Kellogg et al. (2020) examine contestation of algorithmic controls by workers. Virgillito and Ledda (2026) describe challenges to AI-enabled scheduling in nursing. These studies concern the objectives, effects and contestability of managerial control, not direct tests of all other control functions.	A specified function must identify whose purposes it serves and what obligations or interests are relevant. Demonstrating that monitoring or scheduling achieves a managerial objective does not itself establish that the objective, procedures or consequences are justified.

Note: Representative sources illustrate different methods, settings and inferential reach. Connections across sources support the theoretical interpretation rather than a pooled effect estimate.

4.4 Renew: Maintaining expertise and supporting activities

A control that performs satisfactorily today may depend on capabilities whose continued supply is not assured. Renew addresses the activities through which information, professional judgment, software quality and other prerequisites of functional performance are reproduced. The central distinction is between possessing a capability now and sustaining the opportunities that make it available later. Experienced reviewers can resolve unusual payments because they have encountered and discussed related cases. Their present decisions do not demonstrate that employees who replace them will have received comparable opportunities to develop expertise.

The learning mechanism matters. Assigning routine cases to AI can remove some tasks through which employees once learned to discriminate ordinary errors from unusual situations. When those employees are later asked to evaluate complex exceptions, they may have fewer examples to draw upon and less feedback on earlier judgments. Bainbridge (1983) describes the longstanding irony of automation that reserves difficult interventions for people while reducing their practice. Lovett (2026) and Shoaib et al. (2026) develop related concerns about sustained human judgment and the ability to detect and correct subsequent errors. Ertem (2026) examines how changing task assignments affect learning opportunities and how other tasks can supply replacement experience. These works provide plausible mechanisms for delayed effects, not proof that every deployment will erode expertise.

Consider two organizations with equally capable reviewers at the time they begin using AI for payments. In the first, routine cases disappear from employees' work, and the few exceptions that remain arrive without sufficient explanations to support learning. In the second, employees continue to examine a deliberately selected range of cases and receive structured feedback on their judgments. Current exception-resolution rates might initially be similar, yet the organizations differ in their capacity to maintain that function through turnover or changes in case complexity. The prediction follows from the availability of relevant practice and feedback, not simply from the proportion of tasks assigned to AI. Removing repetitive work could even improve learning if the remaining work offers more informative challenges and appropriate support.

Productivity and reported vigilance are distinct from acquired capability. Frieri and Orelli (2026) report productivity in a public-administration Copilot case while treating longer-term learning effects as implications. Naidoo (2026) follows industrial customers over 24 months, finding that initial over-reliance predicts later self-reported complacency and that reported governance safeguards weaken this association. The participants monitor supplier service recovery; they are not organizational AI reviewers, and expertise is not independently assessed. This evidence makes a temporal account of disengagement plausible without establishing the proposed loss of reviewer judgment. Gräfe et al. (2024) and Liu et al. (2027) further identify compensating work and hidden verification effort. Apparent workflow efficiency can therefore coexist with substantial human work whose contribution remains to be assessed.

Renew also extends beyond employees' expertise. AI systems, authorization rules and referral procedures require continuing evaluation as their underlying work changes. Nagbøl et al. (2024) examine evaluation and maintenance at the Danish Business Authority, including weekly quality discussions and human validation alongside resource-constrained formal evaluation. Asatiani et al. (2021) describe cooperation, retraining and suspending models when needed. These studies concern related activity in the same setting and have overlapping authors, so they illustrate connected practices rather than independent replication. They nevertheless show why the organizational work sustaining a technical control cannot be reduced to the existence of the software artifact. Earlier automation research supplies a different warning: Rinta-Kahila et al. (2022) document how removing manual accuracy checks changed verification in Australia's Robodebt scheme. That setting is not equivalent to contemporary agentic AI, but it makes the organizational consequences of withdrawing supporting verification work visible.

An especially consequential relationship arises when an activity that performs one function also supplies a prerequisite for another. In the payment workflow, automatic classification may improve the consistency of routine processing while changing the number and character of cases sent to human reviewers. More referrals can consume the time required for careful resolution; fewer referrals can remove practice that reviewers previously used to develop judgment. These are distinct possible pathways.

A more sensitive software safeguard could increase referrals, reducing the attention available per case even if reviewers' underlying expertise were unchanged. A more selective safeguard could instead leave enough time for difficult cases but reduce the variety of decisions less experienced colleagues see. This difference matters because overloading an existing reviewer and failing to develop future reviewers call for different remedies. The outcomes may occur at different times, so a current performance measure can conceal a change in the conditions of later performance. Improving authorization does not mechanically cause deterioration in exception resolution. To establish a dependency, a study would have to show the connecting activity—changed case exposure, information, review time or feedback—and its effect on the second function. An unrelated training cut or change in staffing could otherwise explain an observed decline.

The same reasoning identifies opportunities for compensation. If another activity supplies a prerequisite lost through delegation, the organization may preserve the affected function without restoring the original task. Simulation, selected manual cases, mentoring, structured feedback and independent checking can replace some opportunities for learning or verification. The effectiveness of compensation is not established by showing that a training session occurred or a second check was created. Investigators need evidence that it preserves the required capability and that the function continues to be performed. A replacement can be incomplete:

simulation may sustain recognition of known cases while offering little preparation for novel situations. Compensation should therefore be assessed against the particular prerequisite it is meant to supply.

Different substitutes may also create new dependencies. An external service that validates transactions can compensate for an internal check, but the adopting organization may then need reliable access to the service's audit evidence and a means of obtaining corrections. A second team might provide independent review at the cost of longer decision times. Compensation preserves a function only when the substitute supplies the relevant capability under the conditions that matter to that function; shifting the activity elsewhere does not eliminate the need to examine its performance.

Social and professional controls present a related question. Ouchi (1979) and Kirsch (1996) explain the importance of shared expectations and informal regulation. AI delegation can reduce occasions when colleagues jointly interpret difficult decisions, while other working arrangements may preserve or even improve those discussions. A technical test will rarely reveal changes in these conditions. Renew thus proposes that later functional performance depends partly on the continued supply of its prerequisites, not merely their present availability. Changes that remove a supporting activity should increase vulnerability when no adequate substitute exists, but deterioration may not occur within the period examined if existing capability remains sufficient for the work or a compensating arrangement sustains it. This is the temporal relationship that a snapshot of current operational adequacy cannot establish.

4.5 Adapt: Recognizing deficiencies and obtaining corrections

Renew helps maintain existing capabilities, but effective control can also require changing an arrangement whose requirements or operating conditions have altered. Adapt concerns the ability to recognize such a mismatch, obtain appropriately authorized action and verify whether the resulting arrangement restores the specified function. A mismatch might be a missing approval requirement, a rule applied to the wrong class of transactions, insufficient reviewer capacity or a new task beyond the expertise of the assigned employees. Recognizing the gap matters, but awareness alone does not change the conditions that produced it.

Responsibility and corrective authority are frequently separated. An employee might discover that a software threshold is outdated, a manager might agree, and an external provider might retain the only permission to revise the code or its configuration. Baird and Maruping (2021) explain the continuing reassessment and coordination required by delegation. Grote et al. (2026) connect accountability to the practical means of exercising control; Yan et al. (2026) examine supervision where verifiability and control over implementation vary. These accounts support a distinction between naming someone accountable and ensuring there is a viable route from evidence of a problem to an effective intervention. The route may cross organizational boundaries and need not give the frontline employee unrestricted authority to change software.

Changes to repeated decisions can require different authority from corrections to individual outputs. Kumar and Suthar (2026) locate discretion in software data, configurations and thresholds. AlQershi et al. (2026), using a survey of 423 professionals and 24 leadership interviews, describe accountability under constrained discretion. Their qualitative accounts concern responsibility without equivalent control; self-reports do not establish actual override behaviour. These sources sharpen a chain of requirements: a deficiency must become observable, be interpreted, reach an authorized actor, permit a feasible intervention and be followed by performance assessment. Correcting one payment does not necessarily change the rule producing similar errors.

Organizational governance and technical-implementation research identifies practices around this chain. Safaei et al. (2026) discuss monitoring and certification during execution, while Ling et al. (2026) examine how developers review, test and revise coding agents' proposed pull requests. These studies concern different settings and do not jointly demonstrate a general improvement in recovery. They show where observations, acceptance decisions and revision practices occur. An escalation record or approved pull request provides evidence that an activity took place; it does not establish that the intended control function subsequently worked. The same distinction applies when a provider reports that a problematic rule has been updated.

Authorized revision entails a tradeoff rather than an argument for unlimited local discretion. Restricting software changes to designated specialists may prevent staff from bypassing safeguards or introducing poorly tested rules. The same restriction may delay correction when local employees see a problem but lack access to those specialists. Distributing authority can improve responsiveness while creating vulnerabilities if changes are made without sufficient knowledge or review. The proposed condition for Adapt is therefore a workable and sufficiently timely route to an authorized, appropriate intervention.

This route can be investigated rather than presumed. Researchers can ask who first observed the problem, whether the evidence was understandable to the receiving actor, what permission was missing, how long escalation took and whether the proposed correction addressed the requirement. A delayed response may reflect an authority gap, but it could instead reflect legitimate testing, disagreement over the proper rule or uncertainty about consequences. Inferences about failed adaptation require attention to those alternatives. Its adequacy depends on the consequence of the error, available time, technical dependencies and the feasibility of reversing an action.

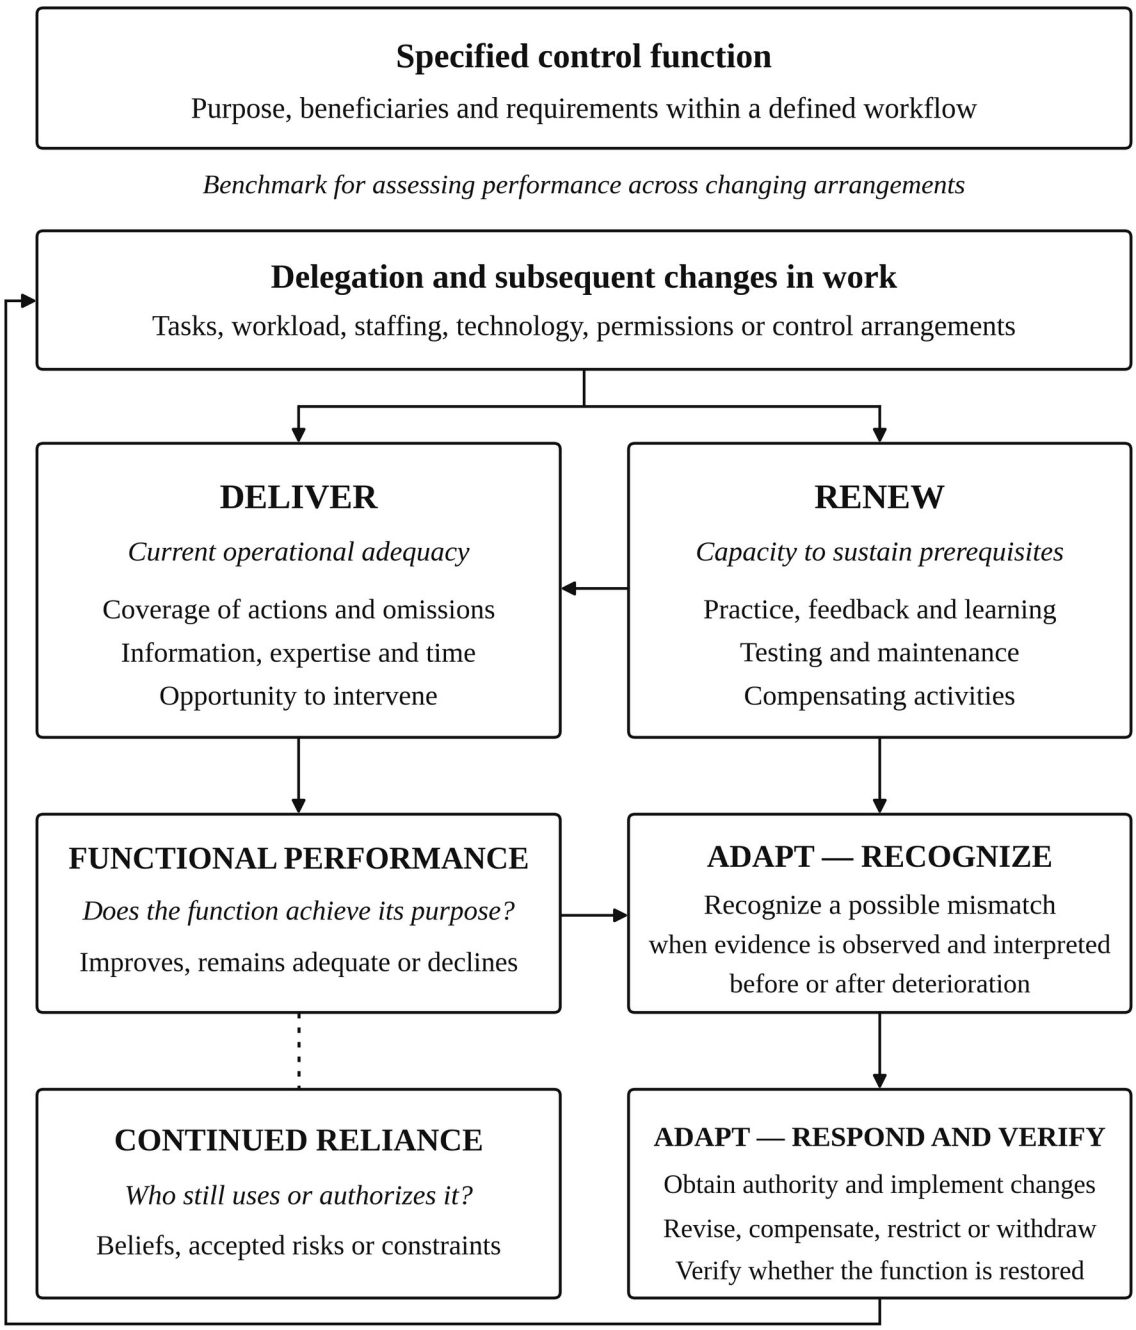
Adaptation can involve additional review, changed permissions, restored human tasks or withdrawal of delegation. Chen et al. (2025) describe professionals contesting AI-related prescription rejections and negotiating rules. Kuziemski and Misuraca (2020) trace complaints, external audit and constitutional review preceding withdrawal of Poland's profiling tool in 2019. This is a correction route involving actors outside the deploying organization; it does not measure whether the intended service function recovered afterward. The case strengthens the account of accessible external authority while preserving the distinction between an intervention and verified recovery. What matters is whether the response addresses the deficiency and subsequent outcomes meet the specified requirement.

Adapt and Renew interact without collapsing into one requirement. Maintenance discussions may expose an emerging coverage problem before it affects performance; changing a safeguard may create new work for reviewers and a new need for training. A correction that improves authorization but overwhelms exception resolution can restore one function and damage another. Equally, strong renewal practices cannot maintain effective control if the specified purpose changes and the unchanged arrangement is no longer suitable. The proposed relationship is that recovery requires recognition plus a viable route to appropriately authorized change and verification. Evidence of the capacity to undertake that sequence should not be confused with evidence that control remained effective throughout a period of failure.

4.6 Functional trajectories, reliance and theoretical contribution

Figure 3 integrates the proposed relationships without depicting Deliver, Renew and Adapt as obligatory successive stages. An arrangement may be delivering its function while employees maintain skills, or may be adapting a deficient control while other functions continue operating. Delegation can alter current prerequisites and the activities sustaining them; maintenance can reveal shortcomings; and an authorized correction can change the conditions for both current performance and future renewal. These links are conditional and can differ across functions within the same workflow. In particular, recognition is not the automatic consequence of functional deterioration: relevant evidence must be generated, interpreted and communicated before anyone can act upon it.

Fig. 3 Proposed relationships explaining dynamic control fit



Note: Corrective responses can alter work arrangements or supporting conditions. Solid arrows indicate proposed conditional influences. Dotted link indicates separate assessment, not causation.

Table 5. Conditional propositions for control after AI delegation

Numbered proposition	Expected relationship following delegation	Evidence needed to evaluate the relationship
P1 · Deliver: coverage	When a delegated safeguard covers more actions and omissions relevant to a specified function, present control will improve if the information, expertise and time needed for intervention remain adequate.	Compare observed authorization results, including omitted actions, before and after the change. A different binding constraint may limit improvement.
P2 · Deliver: interaction	A technical check that changes the number or complexity of referred exceptions can weaken human resolution when the resulting workload exceeds available reviewer capacity.	Follow referrals through the workflow to actual resolution, measuring urgency, case mix and time available for intervention.
P3 · Renew: erosion	When delegation reduces task practice and feedback without compensation, later reviewer judgment will be more likely to decline even if present performance remains adequate.	Track case exposure, subsequent independent judgment, turnover and task demands, with sufficient follow-up to observe potential erosion.
P4 · Renew: compensation	An arrangement can retain its specified function without retaining its original supporting activity when substitute work reproduces the needed capability.	Compare cases with and without supervised practice, feedback, validation or maintenance; test capability and functional outcomes.
P5 · Adapt: authority gap	Recognized control deficiencies will persist longer where those identifying them lack a timely route to an actor who can authorize and implement a proportionate correction.	Trace discovery, interpretation, escalation, formal authority, response time and alternative explanations such as safety testing.
P6 · Adapt: recovery	Corrective intervention restores a deficient function only when subsequent evidence demonstrates that the same specified requirement is again met.	Compare post-change results with the original functional benchmark and assess spillovers to related controls.
P7 · Reliance versus performance	Continued reliance may diverge from assessed performance when actors misunderstand control adequacy, consciously accept limitations or lack feasible alternatives.	Compare contemporaneous beliefs, verified external expectations, constraints, authorized continued use and measured performance.

Note: These propositions synthesize mechanisms across the source streams. They state relationships for future evaluation, not pooled empirical effect estimates.

Propositions 1–7 translate the theoretical explanation into comparisons that can be challenged empirically (Table 5). P1–P2 connect technical coverage with current human intervention; P3–P4 distinguish later capability erosion from compensated continuity; P5–P6 locate recovery in the sequence from recognized mismatch to authorized action and verified outcomes; P7 treats continued reliance as a separate outcome. Coverage, reviewer capacity and revision authority remain different potential constraints, so a remedy must address the weakness actually affecting the specified function.

Continued reliance is a separate outcome from control performance. In ethnographic work at three advertising agencies, Tiwari and Jain (2026) find that practitioners' interpretations of client expectations can limit acted-upon reliance even when they trust system performance; those expectations were not verified with clients. The converse also requires explanation: inadequate controls may remain in use through mistaken confidence, accepted limitations or constrained alternatives. We reserve illusory control persistence for the first case, requiring inadequate performance and contemporaneous beliefs among those sustaining use. Bromley and Powell's (2012) means–ends decoupling distinguishes prescribed activities from achieved results. Dynamic

control fit connects that distinction with changing prerequisites, their supporting work and access to correction. Neither continued use nor restricted reliance establishes actual functional adequacy.

Functional assessment must also specify whose purpose is being pursued. In algorithmic management, AI itself may assign, monitor or evaluate employees rather than be the object of managerial oversight (Schafheitle et al. 2020; Kellogg et al. 2020; Hillebrand et al. 2025; Leavitt et al. 2025). More consistent monitoring may achieve a manager's objective while increasing surveillance or reducing employees' opportunities to contest a decision. Virgillito and Ledda (2026) describe contestation of AI-enabled scheduling by a nursing council; Adams et al. (2026) examine access to AI-generated feedback in police agencies without testing replacement of manual review. These studies show why access to information, voice and the distribution of consequences matter. Whether a control achieves its specified function is an empirical question; whether that function is legitimate or its burdens are acceptable is a separate evaluative question. The two should be investigated together without equating them.

This boundary strengthens the case for functional assessment. Specifying a function identifies the managerial objective under assessment, whose interests it serves and the evidence needed to determine whether it is achieved. Critics may then evaluate the legitimacy of the function and compare alternatives without treating a successful technical check as proof of broader organizational acceptability. Functional clarity is a condition for a sharper normative discussion, not an endorsement of every stated objective.

Dynamic control fit connects established insights about knowledge, measurement and authority through a common functional benchmark. Organizational control theory identifies prerequisites of effective control; AI governance research examines responsibilities and policies; and studies of technical checks and human oversight investigate particular mechanisms. The synthesis explains how these mechanisms affect the work that reproduces their own prerequisites and those of related controls. Its contribution lies in the conditional relationships between current performance, the continuing supply of capability, compensation for displaced activities and feasible correction. These relationships account for divergence between arrangements with similar present performance and for recovery through arrangements different from those originally used.

The explanation is deliberately asymmetric across time. To judge present adequacy, researchers need observed outcomes against a defined requirement. To assess renewal capacity, they must identify the activities that continue to produce the needed expertise or technical reliability and examine whether those activities remain sufficient when work changes. To assess adaptation capacity, they must establish whether the organization can recognize an inadequacy and obtain an effective response. Evidence about renewal and adaptation helps explain the likely durability or recoverability of a function, but cannot convert an observed failure into evidence that the control was continuously effective. A trajectory can contain a period of real loss followed by successful recovery; it can also contain apparently adequate results while prerequisites are being depleted.

Helal et al. (2026) provide a close comparison. Their delegated agency theory links specification, capability-context fit, governance visibility and recovery capacity to perceived delegation fidelity and subsequent evaluations and reliance. Dynamic control fit addresses a different explanatory outcome: the trajectory of a specified control function assessed against its requirements. Consider two payment workflows that currently resolve exceptional transactions equally well. In one, delegation removes routine case experience without replacement; in the other, supervised case review continues to develop the judgment needed for exceptions. Dynamic control fit predicts divergence if that experience is necessary and no other activity compensates for its loss, even while current performance and perceived fidelity remain similar. Following authorization and exception resolution separately also reveals whether stronger enforcement shifts demands onto reviewers faster than their capability develops. The proposed contribution is thus an explanation of how delegation changes the future supply and distribution of control capability, with compensation established through its contribution to performance.

The empirical challenge is to distinguish delegation's effects from concurrent changes in staffing, requirements, complexity or incentives. Comparisons must establish whether a change in supporting activity precedes and helps explain subsequent capability and performance.

For example, the proposed learning relationship would be challenged if a sustained reduction in relevant case exposure and feedback left later reviewer judgment unchanged, despite comparable task demands, no observed replacement practice and sufficient time for the predicted effect to emerge. Capability and performance would need to be assessed independently of the activities expected to sustain them. Evidence of an alternative source of learning could qualify the relationship, but such a substitute must be identified and examined rather than presumed whenever the prediction fails. Section 5 translates the proposed relationships into comparative questions about deterioration, compensation and recovery.

5 Testing dynamic control fit

Dynamic control fit predicts different trajectories for arrangements that appear equally effective at one point in time. Testing that explanation requires specifying the control function, the actions and omissions it should address, the people or interests it serves, and a defensible measure of its actual performance. A study of payments, for example, should assess authorization, referrals and resolution of exceptional transactions separately.

Comparisons should retain the same functional benchmark even where an organization replaces the original arrangement, while recording any deliberate changes to requirements. Table 6 sets out feasible observations and alternative explanations for testing the propositions.

Table 6. Research questions and comparative evidence for testing dynamic control fit

Research question and analytical focus	Comparative design and relevant observations	Competing explanation or qualifying finding
DELIVER — Functional coverage Does delegation preserve or improve a specified control function?	Compare requirements, relevant actions and omissions, and independently assessed results before and after a change in delegation. Evaluate enforcement of permitted and prohibited actions separately from whether required actions were omitted. Record replacement controls and changes in the function being assessed.	An earlier human control was ineffective; a replacement supplies equivalent coverage; requirements changed deliberately; or an observed improvement comes from a wider organizational change rather than delegation.
DELIVER — Human review When can technical referral result in effective resolution?	Measure which cases are referred, their urgency and complexity, reviewer information and expertise, available time, opportunities to intervene, and the correctness and timeliness of final decisions. Distinguish referral performance from the resolution of the cases referred.	Better selection of cases reduces unnecessary work; reviewer shortfalls reflect unrelated workload changes; or apparently fewer interventions result from undetected problems rather than improved processing.
RENEW — Sustaining prerequisites Does reduced case exposure precede a loss of assessed reviewer capability, rather than only reported vigilance?	Follow exposure, practice, feedback, turnover and maintenance before subsequent independent judgment assessments. Compare similar task demands with and without replacement practice. Measure vigilance separately from expertise; customers' reported complacency cannot substitute for reviewer performance.	Capability remains sufficient despite reduced exposure and no observed replacement practice, with comparable tasks and adequate follow-up; alternatively, deterioration reflects staffing or workload changes. A compensating activity must be observed rather than inferred from preserved performance.
DELIVER / RENEW — Dependencies and compensation Can strengthening one function alter another?	Trace changes in routing or enforcement through case volume, exposure to relevant judgments, reviewer workload, feedback or technical maintenance, then assess a second function. Compare settings with and without compensating activities, where feasible.	Both functions respond independently to a broader change, or replacement activities maintain the relevant prerequisite despite changes to the original workflow.

Research question and analytical focus	Comparative design and relevant observations	Competing explanation or qualifying finding
ADAPT — Recognition and corrective authority Why do known control deficiencies remain unresolved?	Follow identifiable signals of mismatch through interpretation, escalation, access to decision makers, authorization and implementation. Use incident records, permissions, communications and change logs, while distinguishing the recorded discovery date from the uncertain onset of the problem.	Delay arises from technical complexity or uncertainty rather than inaccessible authority; a different actor can obtain a timely correction; or greater local discretion introduces unsafe revisions.
ADAPT — Verified recovery Does the specified function recover after revision, withdrawal or replacement?	Follow local, provider and external correction routes through implementation, then assess the same deficient function against unchanged performance criteria. Examine effects on other functions. Distinguish intervention, including withdrawal, from independently verified recovery.	Performance improves because conditions changed independently; the original requirement remains unmet; or a correction restores one function while impairing another.
CROSS-CUTTING — Performance and reliance When does reliance follow performance, external expectations or constraints?	Compare measured performance with contemporaneous beliefs, attributed and verified client expectations, resource decisions and continued-use authorizations. Identify whose decisions sustain or restrict reliance and what those actors understood when acting.	Actors accept a known deficiency or lack an alternative; client expectations restrict use despite performance trust. Neither continued use with poor performance nor limited reliance establishes mistaken confidence without direct belief evidence.
BOUNDARIES — Technology and affected groups Where does the explanation apply, and to whose control objectives?	Compare similar functions across AI decision support, connected agents and AI-directed managerial work. Specify task complexity, reversibility, delegation rights, stakeholder impacts and opportunities to contest decisions alongside functional outcomes.	Differences reflect the task, organizational setting or objective rather than the delegation technology. A control can achieve its stated function while imposing costs or serving objectives that affected participants reasonably contest.

The unresolved learning relationship requires observing case exposure and feedback before independently assessed changes in reviewer judgment. Naidoo's (2026) two-wave findings concern customers' reported complacency, leaving reviewer capability untested. Comparable workflows could be followed through delegation and staff turnover, recording case difficulty, practice, maintenance and replacement assignments. Stable judgment after reduced exposure would challenge predicted erosion when task demands, time horizon and unobserved compensation have been adequately examined. Turnover or changed case complexity could otherwise explain deterioration. Such comparisons distinguish changes in motivation or vigilance from loss of the knowledge needed to resolve difficult cases.

For Adapt, the question extends beyond whether correction is possible to whether a function recovers after revision, restriction, withdrawal or replacement. Kuziemski and Misuraca (2020) document an external correction route without measuring subsequent recovery. Incident reports, permissions, provider communications and change logs can trace signals through escalation and implementation; consistent performance criteria must establish the result. For reliance, researchers should compare measured performance with contemporaneous beliefs, client expectations and continued-use decisions. Tiwari and Jain (2026) identify externally shaped reliance, but neither that finding nor continued use establishes mistaken confidence. The first recorded incident need not mark the onset of failure.

Comparisons should also vary task reversibility, technical architecture and the placement of authority, including settings in which AI exercises managerial control over employees. A control can achieve its stated function while creating unjustified surveillance or transferring burdens onto affected people. These comparisons would establish the conditions under which the proposed relationships operate, fail or require qualification. These comparisons can test the framework across organizational and technical settings.

Evidence of dependencies would be especially valuable where an intervention improves one function and appears to weaken another. Case-routing records can show whether reviewers encounter a narrower range of decisions after automation, while subsequent judgment tests can establish whether this experience matters for their ability to resolve exceptions. Researchers should distinguish that mechanism from a common external cause affecting both outcomes. Qualitative process tracing can complement performance measures by documenting who noticed a deficiency, how it was interpreted, and why the needed correction did or did not follow. These designs preserve a difference between explanations that are theoretically credible and those supported by a particular empirical trajectory.

6 Assessing control in organizational practice

Dynamic control fit offers managers and auditors a way to examine what governs AI-enabled work, rather than simply confirming the existence of policies or safeguards. The assessment should begin by naming a particular function and the evidence that would establish its performance. In a payment workflow, preventing unauthorized transactions and resolving exceptional cases are distinct purposes. A technically reliable threshold check may succeed at the first while leaving gaps in approvals that the system never requests, or sending more cases to employees than they can resolve before execution.

For Deliver, managers should establish coverage, meaningful information, reviewer expertise and a timely opportunity to intervene. The outcome is achievement of the specified function, beyond success on a software test. For Renew, they should identify how practice, feedback and maintenance reproduce these prerequisites. Reducing manual decisions can improve throughput while changing employees' learning opportunities. Structured practice and revised assignments are possible responses, but evidence of later capability is needed to establish compensation. Reported confidence or vigilance cannot substitute for that assessment, and the review does not establish a generally effective training prescription. Software and provider-supported rules likewise require evaluation against current requirements.

For Adapt, managers should establish who can identify a failure, who receives and evaluates the evidence, and who has permission and resources to change the arrangement. A local reviewer may recognize an obsolete rule while the authority to revise it belongs to a provider. Remedies include software correction, revised referral processes, additional review and temporary restrictions on AI permissions. A change request is evidence that action occurred; its effectiveness must be checked against subsequent outcomes and possible effects on other functions. Missing coverage, insufficient reviewer capacity and inaccessible revision authority require different remedies.

These assessments should distinguish performance from reliance. Mistaken confidence requires a different response from known limitations or an unavailable alternative. Client expectations can also restrict use even when staff trust a system, so managers should record whose expectations influence delegation and verify them directly. They should also identify whose objectives a control serves, who bears checking and correction work, and whether affected people can challenge decisions.

There can be legitimate organizational reasons to keep an imperfect arrangement temporarily, such as the absence of a workable replacement or the risk that an immediate change would cause a more serious failure. This is not the same as mistakenly believing it is adequate. Managers should record what is known about the limitation, what temporary constraints reduce its consequences, who authorized continued use and when the function will be reassessed. Such documentation cannot substitute for performance, but it helps distinguish conscious acceptance of risk from an inability to detect deterioration. Deliver, Renew and Adapt direct repeated assessments after changes in staffing, tasks, technology or requirements. They are not a prescribed sequence or a validated checklist score.

An organization can use these distinctions to allocate its audit effort more effectively. A known and stable approval rule might require periodic tests of edge cases and omissions; an arrangement with many exceptional referrals may require a study of reviewer workloads and judgment; and a system maintained externally may

require documented escalation and testing rights. The evidence needed depends on what can go wrong in that particular control function and which conditions appear most vulnerable.

For example, a manager confronted with a surge in referrals should not immediately infer that the software's enforcement rules are defective. The safeguard might now catch more genuine exceptions, or it might refer routine cases unnecessarily. Employees might also have lost capacity because staffing changed. Reviewing the actual cases and the ability to resolve them identifies whether the failure lies in coverage, workload or judgment, and therefore which response is proportionate. Similarly, a quarterly provider report may describe completed maintenance without establishing that the current approval requirement is correctly implemented. Independent functional testing can make the difference visible.

7 Explaining sustained control as work is delegated to AI

Organizations can retain familiar policies and supervisory roles after delegating work to AI while changing the conditions that make those controls effective. A newly implemented technical check can improve authorization, yet the same change may shift verification work, limit human learning or place corrective authority beyond the people who recognize a problem. The relevant theoretical question is therefore not whether an arrangement remains in place, but how a specified control function performs as the organization of work changes.

Dynamic control fit addresses that question by connecting three requirements. Deliver concerns current coverage, judgment and opportunities for intervention. Renew concerns the practice, maintenance, feedback and compensating activities on which later performance depends. Adapt concerns recognizing an emerging mismatch, obtaining appropriately authorized action and establishing whether the function has actually been restored. These relationships explain how an arrangement can work now while becoming vulnerable later, why one improved control can weaken another, and how a weakened function can be sustained or recovered through a different arrangement. Recoverability does not establish uninterrupted effectiveness.

The synthesis connects functional performance with the activities supplying its prerequisites and the feasibility of verified correction. Continued reliance is a separate outcome: beliefs, external expectations and practical constraints can sustain use despite declining effectiveness or limit use despite adequate performance. The resulting explanation yields testable comparisons between arrangements that perform similarly today but differ in learning opportunities, compensating activities or routes to authorized correction. Following the same function over time allows researchers to distinguish these mechanisms from changes in task demands, staffing and external expectations.

These comparisons make changes in control arrangements theoretically consequential. Preserved performance through new supporting activities differs from deterioration behind an unchanged formal control, just as verified recovery differs from uninterrupted effectiveness. The functional benchmark permits both to be assessed without assuming that retaining or replacing a particular mechanism determines the outcome.

Specifying a function also permits scrutiny of whether its objective is justified, whose interests it serves and what consequences it imposes, particularly when AI directs managerial control over workers.

References

Adams IT, McLean K, Alpert GP (2026) Improving police behavior through artificial intelligence: Pre-registered experimental results in two large US agencies. Criminology 64:502–516. https://doi.org/10.1111/1745-9125.70028

AlQershi N, Thurasamy R, Alzoraiki M, Ali GA, Emam ASM (2026) Agentic leadership in the algorithmic age: leveraging AI and ethical intelligence for crisis-resilient organizations. Leadership & Organization Development Journal 47(6):1205–1223. https://doi.org/10.1108/LODJ-06-2025-0512

Ante L (2026) Specification cost: Relocating the binding constraint in delegation to AI agents with externally supplied objectives. Technology in Society 88:103503. https://doi.org/10.1016/j.techsoc.2026.103503

Arora C, Vogelsang A, Sharma A (2026) Specifying the delegated-autonomy boundary: Requirements engineering for agentic AI. arXiv preprint arXiv:2607.17225. https://doi.org/10.48550/arXiv.2607.17225

Asatiani A, Malo P, Nagbøl PR, Penttinen E, Rinta-Kahila T, Salovaara A (2021) Sociotechnical envelopment of artificial intelligence: An approach to organizational deployment of inscrutable artificial intelligence systems. Journal of the Association for Information Systems 22(2):325–352. https://doi.org/10.17705/1jais.00664

Bainbridge L (1983) Ironies of automation. Automatica 19(6):775–779. https://doi.org/10.1016/0005-1098(83)90046-8

Baird A, Maruping LM (2021) The next generation of research on IS use: A theoretical framework of delegation to and from agentic IS artifacts. MIS Quarterly 45(1):315–341. https://doi.org/10.25300/misq/2021/15882

Basu P (2026) Governing agentic performance management: A sociotechnical framework for organisational effectiveness. International Journal of Organizational Analysis, 1–18. https://doi.org/10.1108/ijoa-01-2026-6408

Birkstedt T, Minkkinen M, Tandon A, Mäntymäki M (2023) AI governance: Themes, knowledge gaps and future agendas. Internet Research 33(7):133–167. https://doi.org/10.1108/intr-01-2022-0042

Block J, Brändle L, Limbach P, Röth T, Sabel CA (2025) Half dead or more alive than ever? New, old, and (increasingly) redundant literature review formats. Management Review Quarterly 75:993–1008. https://doi.org/10.1007/s11301-025-00513-9

Bromley P, Powell WW (2012) From smoke and mirrors to walking the talk: Decoupling in the contemporary world. Academy of Management Annals 6(1):483–530. https://doi.org/10.1080/19416520.2012.684462

Cardinal LB, Sitkin SB, Long CP (2004) Balancing and rebalancing in the creation and evolution of organizational control. Organization Science 15(4):411–431. https://doi.org/10.1287/orsc.1040.0084

Cardinal LB, Kreutzer M, Miller CC (2017) An aspirational view of organizational control research: Re-invigorating empirical work to better meet the challenges of 21st century organizations. Academy of Management Annals 11(2):559–592. https://doi.org/10.5465/annals.2014.0086

Campbell R (2026) Agentic shadow infrastructure: How AI supply-chain drift creates unmanaged enterprise infrastructure. Computers 15(8):510. https://doi.org/10.3390/computers15080510

Chen Q, Wang H, Ma C, Ru P (2025) How professionals respond to disruptive effects of artificial intelligence on their jurisdiction: The role of interactive governance. Social Science & Medicine 386:118626. https://doi.org/10.1016/j.socscimed.2025.118626

Cuneo J, Chun D, Khanna G (2026) AI-GRACE: A use-case operationalization framework for agentic AI: From organizational objectives and obligations to deployment capabilities and architecture. arXiv preprint arXiv:2609.21192. https://doi.org/10.48550/arXiv.2609.21192

Elsbach KD, van Knippenberg D (2020) Creating high-impact literature reviews: An argument for 'integrative reviews'. Journal of Management Studies 57(6):1277–1289. https://doi.org/10.1111/joms.12581

Ertem O (2026) AI-Mediated Apprenticeship: Rethinking Follower Development in Servant Leadership Under GenAI Delegation. Administrative Sciences 16:505. https://doi.org/10.3390/admsci16100505

Frieri FR, Orelli RL (2026) Governing generative AI in public administration: Productivity, work redesign and reflective human–AI collaboration. Public Money & Management, advance online publication. https://doi.org/10.1080/09540962.2026.2703173

Gräfe P, Marienfeldt J, Wehmeier LM, Kuhlmann S (2024) Changing tasks and changing public servants? The digitalisation and automation of public administrative work. Information Polity, 13837605241289773. https://doi.org/10.1177/13837605241289773

Grote G, Parker SK, Crowston K (2026) Taming artificial intelligence: A theory of control–accountability alignment among AI developers and users. Academy of Management Review 51(2):278–299. https://doi.org/10.5465/amr.2023.0117

Helal MYI, Dwivedi YK, Ismagilova E, Chae I (2026) Reinterpreting agency theory in the era of artificial intelligence: conceptualizing delegated agency. AI & Society 41:7063–7074. https://doi.org/10.1007/s00146-026-03034-5

Hillebrand L, Raisch S, Schad J (2025) Managing with artificial intelligence: An integrative framework. Academy of Management Annals 19(1):343–375. https://doi.org/10.5465/annals.2022.0072

Jeppe A, Bree T, Karger E, Ahlemann F, Proff H (2026) Artificial intelligence as manager: confronting implementation barriers and shaping future research. Management Review Quarterly. https://doi.org/10.1007/s11301-026-00622-z

Jørgensen BN, Ma ZG (2026) Governing agentic AI in enterprise workflows: A bounded-autonomy framework for delegated authority and controlled execution. Information 17(9):923. https://doi.org/10.3390/info17090923

Kathuria A, Karhade PP, Malik O, Baruri VKP, Konsynski BR (2026) AI governance and the decentralization of technology production: An investigation of AI-based IPA bots. Information Systems Research 37(3):1873–1894. https://doi.org/10.1287/isre.2023.0588

Kellogg KC, Valentine MA, Christin A (2020) Algorithms at work: The new contested terrain of control. Academy of Management Annals 14(1):366–410. https://doi.org/10.5465/annals.2018.0174

Kim H, Glaeser EL, Hillis A, Kominers SD, Luca M (2024) Decision authority and the returns to algorithms. Strategic Management Journal 45(4):619–648. https://doi.org/10.1002/smj.3569

Kirsch LJ (1996) The management of complex tasks in organizations: Controlling the systems development process. Organization Science 7(1):1–21. https://doi.org/10.1287/orsc.7.1.1

Kumar D, Suthar N (2026) Programmes as decision premises in AI-mediated organisations: A social systems theory perspective. Information Technology & People 39(6):2373–2387. https://doi.org/10.1108/ITP-12-2025-2038

Kuziemski M, Misuraca G (2020) AI governance in the public sector: Three tales from the frontiers of automated decision-making in democratic settings. Telecommunications Policy 44:101976. https://doi.org/10.1016/j.telpol.2020.101976

Leavitt K, Barnes CM, Shapiro DL (2025) The role of human managers within algorithmic performance management systems: A process model of employee trust in managers through reflexivity. Academy of Management Review 50(4):745–767. https://doi.org/10.5465/amr.2022.0058

Lebovitz S, Levina N, Lifshitz-Assaf H (2021) Is AI ground truth really true? The dangers of training and evaluating AI tools based on experts' know-what. MIS Quarterly 45(3):1501–1525. https://doi.org/10.25300/MISQ/2021/16564

Liel Y, Zalmanson L (2025) Turning off your better judgment: Algorithmic conformity in artificial intelligence-human collaboration. Journal of Management Information Systems 42(4):1087–1117. https://doi.org/10.1080/07421222.2025.2561390

Ling M, Zou J, Liu X (2026) Negotiating Completion in Agentic Software Work: Acceptance Boundaries and Deferred Requirements. Systems 14:1312. https://doi.org/10.3390/systems14101312

Liu J, Zhong C, Meng F, Shank DB, Wunsch DC (2027) GenAI-enabled human–AI collaboration: An integrative review of local augmentation, hidden work, and governable synergy. International Journal of Information Management 92:103140. https://doi.org/10.1016/j.ijinfomgt.2026.103140

Lovett N (2026) The tragedy of the cognitive commons: How AI could disrupt the regeneration of professional expertise. Human Resource Development Review, 15344843261470602. https://doi.org/10.1177/15344843261470602

Mäntymäki M, Minkkinen M, Birkstedt T, Viljanen M (2022) Defining organizational AI governance. AI and Ethics 2(4):603–609. https://doi.org/10.1007/s43681-022-00143-x

Mökander J, Sheth M, Gersbro-Sundler M, Blomgren P, Floridi L (2022) Challenges and best practices in corporate AI governance: Lessons from the biopharmaceutical industry. Frontiers in Computer Science 4:1068361. https://doi.org/10.3389/fcomp.2022.1068361

Nagbøl PR, Krancher O, Müller O (2024) Challenges and design principles for the evaluation of productive AI systems in the public sector. In Charalabidis, Y., Medaglia, R. and van Noordt, C. (eds), Research Handbook on Public Management and Artificial Intelligence. Edward Elgar Publishing, 245–262. https://doi.org/10.4337/9781802207347.00025

Naidoo V (2026) The dark side of AI-enabled industrial service recovery: The erosion of trust and customer engagement. Industrial Marketing Management 135:205–223. https://doi.org/10.1016/j.indmarman.2026.05.004

Ouchi WG (1979) A conceptual framework for the design of organizational control mechanisms. Management Science 25(9):833–848. https://doi.org/10.1287/mnsc.25.9.833

Pacella M, Papadia G, Giliberti V (2026) Governed agentic process automation: A floor-safety guarantee for compliance-critical LLM routing. Algorithms 19(8):627. https://doi.org/10.3390/a19080627

Papagiannidis E, Mikalef P, Conboy K (2025) Responsible artificial intelligence governance: A review and research framework. The Journal of Strategic Information Systems 34(2):101885. https://doi.org/10.1016/j.jsis.2024.101885

Rinta-Kahila T, Someh I, Gillespie N, Indulska M, Gregor S (2022) Algorithmic decision-making and system destructiveness: A case of automatic debt recovery. European Journal of Information Systems 31(3):313–338. https://doi.org/10.1080/0960085x.2021.1960905

Safaei D, Hein A, Haki K (2026) Logic hybridization in enterprise platforms: Reconfiguring architecture and governance for generative AI. Journal of Information Technology 41(3):274–304. https://doi.org/10.1177/02683962261467413

Schafheitle SD, Weibel A, Ebert IL, Kasper G, Schank C, Leicht-Deobald U (2020) No stone left unturned? Toward a framework for the impact of datafication technologies on organizational control. Academy of Management Discoveries 6(3):455–487. https://doi.org/10.5465/amd.2019.0002

Schulte N, Kanbach DK (2026) Rethinking organizational decision-making: The emerging roles and tasks of generative artificial intelligence. Management Review Quarterly. https://doi.org/10.1007/s11301-026-00611-2

Shoaib S, ul Haq F, Maaitah N, Dhahbi A, Badawy M, Sheikh AA, Masood A (2026) A multilevel framework of human–AI knowledge work and its epistemic consequences for organizational decision-making. Social Sciences & Humanities Open 14:103799. https://doi.org/10.1016/j.ssaho.2026.103799

Tarafdar M, Rets I, Hu Y (2026) Catch me if you can: The dynamic nature of bias in machine learning applications. Information Systems Journal 36(5):710–742. https://doi.org/10.1111/isj.70034

Tiwari N, Jain V (2026) Governing AI across markets: How client-origin expectations shape human-AI collaboration in globally oriented advertising agencies in Asia. Journal of Global Marketing, advance online publication. https://doi.org/10.1080/08911762.2026.2733411

Torraco RJ (2016) Writing integrative literature reviews: Using the past and present to explore the future. Human Resource Development Review 15(4):404–428. https://doi.org/10.1177/1534484316671606

van Esch P, Cui YG (2026) Agentic brand drift: How AI-orchestrated organizations will lose their identity and how to get it back. Business Horizons, advance online publication. https://doi.org/10.1016/j.bushor.2026.05.013

Virgillito D, Ledda C (2026) AI-Mediated Participation and People Sustainability: A Socio-Technical Case Study in Healthcare Shift Scheduling. Systems 14:168. https://doi.org/10.3390/systems14020168

Yan H, Sun Z, Liu S (2026) Managerial altitude: Where should humans stand when AI agents manage AI agents? Business Horizons. https://doi.org/10.1016/j.bushor.2026.09.001

Zheng H, Dong Q, Depena RK, Bhatia JD, Xiao F, Xu P (2026) Separating capability from permission: A governance framework for agentic AI autonomy levels. arXiv preprint arXiv:2607.23438. https://doi.org/10.48550/arXiv.2607.23438

Zwerdling N, Boaz D, Rabinovich E, Uziel G, Amid D, Anaby Tavor A (2025) Towards enforcing company policy adherence in agentic workflows. Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: Industry Track, 595–606. https://doi.org/10.18653/v1/2025.emnlp-industry.41

Statements and declarations

Data availability

Online Resource 1 describes the review procedures and their limits. Online Resource 2 provides final search protocols, full-text decisions, the source inventory, primary analytical assignments and selected evidence appraisals. Full texts remain with their publishers or repositories and are not redistributed.

Supplementary information

Online Resource 1: Search, selection and analytical procedures. Online Resource 2: Source inventory, coding and claim evidence.

Funding

The author declares that no funds, grants or other support were received for this work.

Competing interests

The author declares no relevant financial or non-financial competing interests.

Author contributions

The sole author conceived and designed the review, conducted the literature search and screening, analysed and interpreted the sources, developed the theoretical framework, and prepared and approved the manuscript.

Use of generative AI

Generative-AI tools, including ChatGPT (OpenAI), assisted with selected copy-editing, table preparation and spreadsheet formatting. The author reviewed and revised these outputs and takes responsibility for the manuscript.