ORIGINAL RESEARCH **Open access**

Deliberate AI Non-Adoption as a Strategic Capability for Preserving Critical Expertise, Reversible Commitments, and Organizational Control over Consequential Tasks

Peter Lindberg[1*], Thomas Berger[2], Marta Nowak[3]

Abstract

Artificial intelligence strategy is commonly framed as a problem of adoption speed, deployment scope, or effective human–AI collaboration. That framing leaves unresolved whether organizations may sometimes create strategic value by intentionally withholding AI from selected consequential tasks. This article develops a Prescriptive Design Theory of deliberate AI non-adoption as a bounded organizational capability. The argument integrates research on heterogeneous AI performance, organizational decision structures, expertise development, human–AI delegation, algorithmic control, and reversibility. Deliberate non-adoption is defined as an intentional, task-specific, governable, and reviewable decision to preserve human performance of selected activities while permitting AI use elsewhere. Its strategic value depends on three coupled conditions: continued human task exposure must materially preserve critical expertise; present deployment must otherwise create commitments that are difficult to reverse; and effective organizational control must depend on capabilities that formal override rights alone cannot secure. The theory therefore rejects both generalized AI avoidance and automatic adoption. Deliberate non-adoption is defensible only when its preservation benefits exceed the performance, learning, or coordination opportunities forgone by withholding AI. The resulting framework treats AI strategy as a reversible allocation problem across heterogeneous tasks rather than as a unidirectional transition toward greater automation.

Keywords Artificial intelligence strategy, Deliberate non-adoption, Heterogeneous tasks, Unidirectional transition

*Correspondence:
Peter Lindberg
peter.lindberg@uu.se
[1] Department of Strategy and Innovation, Uppsala University, Uppsala, Sweden
[2] Department of Management, University of Bern, Bern, Switzerland
[3] Department of Management, University of Warsaw, Warsaw, Poland

Introduction

Artificial intelligence has expanded the set of organizational tasks that can be automated, accelerated, or recombined with human judgment. Yet the strategic problem created by this expansion is not adequately described by asking whether organizations should automate or augment. Automation and augmentation are interdependent arrangements whose effects emerge through the organization of work, authority, and human–machine interaction [1]. A system perspective similarly shows that AI becomes consequential through its position within a broader configuration of design, implementation, use, professional practice, and organizational coordination [2]. AI strategy therefore concerns the allocation of tasks, decision rights, learning opportunities, and accountability across a sociotechnical system.

This is an open access article distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0), which permits non-commercial use, distribution, reproduction, and adaptation in any medium or format, provided that appropriate credit is given to the original author(s) and source, a link to the license is provided, and any changes made are indicated. Adapted material must be distributed under the same license.

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

Recent evidence makes that allocation problem more difficult because the effects of AI are sharply heterogeneous. In controlled knowledge-work settings, generative AI can increase speed and quality for tasks that fall within its capability frontier while lowering accuracy when workers rely on it for tasks outside that frontier [3, 4]. Large-scale deployment evidence also indicates that average productivity gains can coexist with substantial variation across workers and task types. These findings weaken strategies built on uniform assumptions about the relative competence of humans and AI. The relevant question is increasingly which activities should be delegated, augmented, retained, monitored, or periodically reallocated.

The governance literature adds a second complication. Responsible AI requires organizational structures, procedures, relational practices, and mechanisms of accountability that extend beyond model performance [5]. Control becomes especially important when adaptive systems affect outcomes that users or developers cannot fully predict or later reconstruct. Control and accountability can become misaligned when actors remain answerable for outcomes while their practical ability to influence those outcomes declines [6]. Organizations may therefore possess formal responsibility without retaining the knowledge, authority, or operational competence necessary to exercise meaningful control.

These developments reveal a strategic issue that adoption-centered accounts rarely isolate. Some organizational capabilities are sustained through continued performance of consequential work. If AI progressively displaces that work, the organization may gain immediate efficiency while weakening its ability to make unaided judgments, recognize failure, contest recommendations, or resume human execution when technological or institutional conditions change. Conversely, preserving human participation indiscriminately can sacrifice measurable AI benefits, slow learning about emerging systems, and protect routines whose strategic value has already diminished.

The theoretical problem is consequently conditional. The question is not whether non-adoption is intrinsically desirable. It is when withholding AI from a bounded class of tasks can preserve capabilities whose future value depends on continued human exercise, while avoiding the costs of generalized technological resistance. Prescriptive Design Theory is appropriate because the issue concerns choices among organizational arrangements under explicit mechanisms, constraints, trade-offs, and failure risks. The article develops design principles for deciding where deliberate non-adoption may be justified, where hybrid allocation is preferable, and where greater AI use remains the stronger arrangement.

Beyond Adoption-Centric AI Strategy

An adoption-centric view implicitly treats greater deployment as the principal strategic trajectory and interprets organizational differences largely through the speed, breadth, or effectiveness with which AI is introduced. That view becomes incomplete once AI affects the internal allocation of decision authority. Organizations can delegate decisions entirely to AI, use sequential human–AI arrangements, aggregate judgments across human and machine actors, or retain selected decisions within human authority [7]. These alternatives are organizational structures with different implications for speed, interpretability, accountability, learning, and error recovery.

Decision rights can also be reapportioned at a finer level than an entire task. A decision process may contain sensing, diagnosis, option generation, evaluation, authorization, execution, and review components, and AI competence need not be equivalent across them. Cognitive reapportionment therefore makes it possible to redistribute authority dynamically between human and computational actors [8]. An organization can expand AI use while preserving human authority over selected elements, or conversely maintain human execution while using AI for analysis, simulation, retrieval, or challenge. Adoption intensity alone obscures these distinctions.

The behavioral acceptance of AI does not resolve the allocation problem. Perceptions of AI agency can create opposing trust effects: greater perceived capability may increase reliance, while the same perception of agency may intensify concerns about betrayal, unpredictability, or loss of control [9]. Acceptance can also arise under constrained choice. Longitudinal research on algorithmically managed platform work shows how repeated, bounded choices can generate

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

consent within a system that continues to structure workers' available options [10]. Observable participation in an AI-mediated arrangement therefore does not establish that the arrangement preserves meaningful organizational discretion.

A strategy centered only on increasing adoption can consequently overlook three forms of path dependence. First, the organization may progressively remove human exposure to activities through which important expertise is developed. Second, workflow and governance systems may become increasingly dependent on AI-mediated routines that are costly to reverse. Third, formal decision authority may remain nominally human while the practical knowledge needed to challenge machine recommendations deteriorates. These processes can develop gradually and may be difficult to identify from near-term performance measures.

This argument does not establish a general case for withholding AI. Adoption may improve performance, spread specialized knowledge, reduce error, and release human capacity for activities in which human judgment contributes more. The strategic deficiency lies in treating greater adoption as a self-validating objective. Once organizations recognize that different task components carry different learning, control, and reversibility consequences, non-use becomes one possible design choice within a broader portfolio of human–AI arrangements.

Deliberate Non-Adoption as Strategic Choice

Deliberate AI non-adoption is defined here as an intentional, bounded, governable, and periodically reviewable decision to withhold a specified form of AI use from a specified class of tasks while permitting AI deployment elsewhere. Four features distinguish it from ordinary non-use. Intentionality requires an articulated organizational reason. Boundedness requires identification of the affected tasks, decisions, or authority domains. Governability requires responsibility for maintaining and enforcing the allocation. Reviewability requires that the arrangement can be reconsidered as technologies, organizational capabilities, or external conditions change.

This definition excludes several superficially similar states. An organization that lacks funds, data, technical infrastructure, regulatory permission, or managerial attention has not necessarily developed a non-adoption capability. Employee resistance alone is also insufficient. Deliberate non-adoption requires a positive design rationale linked to an organizational mechanism whose value can be challenged against alternative arrangements.

One such rationale arises from uncertainty. Prediction systems are strongest where relevant states can be represented and learned from sufficiently informative data, but organizational decisions can also involve novelty, indeterminism, recursive reactions, and forms of Knightian uncertainty that cannot be reduced to estimated probabilities [11]. Where uncertainty is materially irreducible, retaining human involvement may preserve forms of inquiry, interpretation, and adaptive judgment that a predictive system does not replace. This does not imply that humans solve Knightian uncertainty reliably; it means only that prediction quality cannot by itself determine the appropriate allocation of authority.

A second rationale concerns corrective human work. Human-in-the-loop arrangements can evolve beyond symbolic oversight when employees audit, correct, reinterpret, and contextualize algorithmic outputs [12]. Continued human participation can therefore perform an active governance function. Yet this function is contingent on the quality of the expertise being preserved. Research on expert–AI pairings shows that intervention depends on professional knowledge, beliefs about the technology, organizational priorities, and the way experts interpret AI as either an ally or a rival [13]. Human retention without demonstrable task-relevant competence can preserve error as easily as expertise.

Deliberate non-adoption should thus be understood as one allocation option among full delegation, bounded delegation, augmentation, mandatory human review, and human execution supported by non-decisional AI. Its strategic justification must identify a specific mechanism that would be weakened by deployment. Where no such mechanism exists, withholding AI may amount to avoidable inefficiency or institutional inertia.

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

Taken together, these distinctions locate deliberate non-adoption within a wider allocation architecture in which organizations choose where AI may act, where human performance must remain active, and which capabilities must remain available for future reallocation.

The relationships discussed here are shown in **Figure 1**.

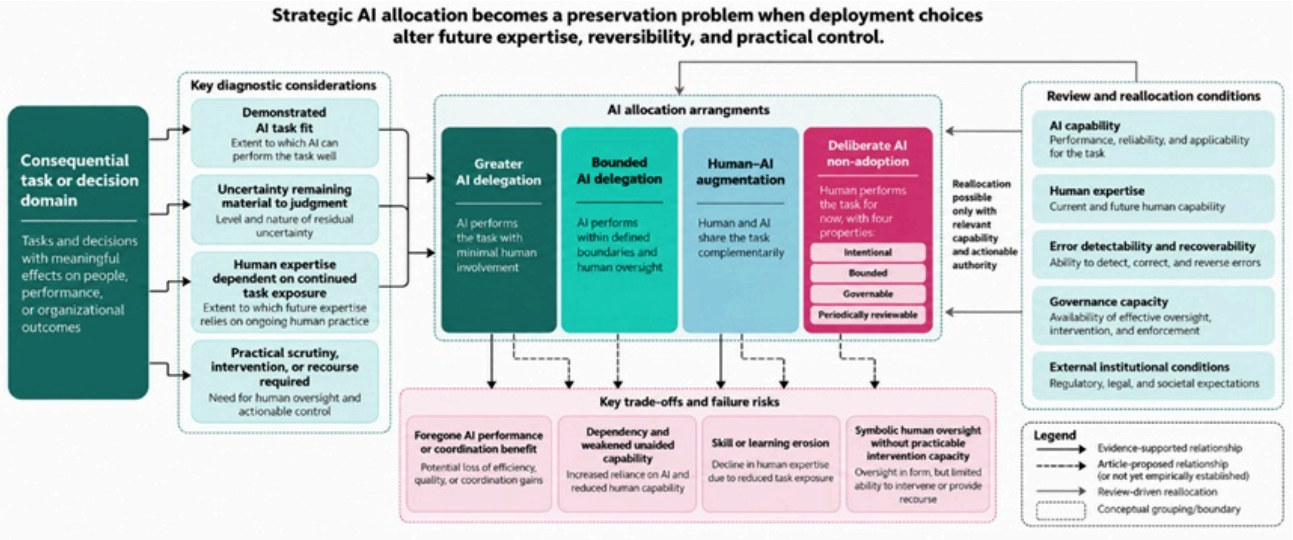

Figure 1. Strategic AI allocation becomes a preservation problem when deployment choices alter future expertise, reversibility, and practical control

Preservation of Critical Expertise

The strongest strategic case for deliberate non-adoption emerges when expertise is sustained through repeated exposure to the very task being considered for delegation. Decision competence is not a static stock that remains available once work is removed. Experimental evidence indicates that reliance on machine-learning decision aids can improve immediate performance while weakening the development of unaided decision skill [14, 15]. The strategic implication is conditional: where future human performance matters and competence decays or fails to develop without practice, current AI assistance can create a capability trade-off that is invisible in contemporaneous productivity measures.

Expertise preservation also operates above the individual level. Organizational learning depends on the range of actions, interpretations, outcomes, and feedback available to members. Substituting machine learning for human decision-making can change the richness of that learning process by altering who acts, what feedback humans observe, and which decision routines remain practiced [16]. An organization may therefore retain employees with nominal domain knowledge while losing the collective routines through which judgment is updated and transmitted.

Professional expertise introduces a further distinction. AI can be interpreted differently by actors occupying different positions within the same profession. Research in radiology shows that front-line professionals, hybrid experts, and field-level actors can understand the same epistemic technology as diminishing, complementing, or enhancing professional agency [17]. Critical expertise is therefore not equivalent to preserving every legacy task. Some tasks can migrate to AI while human roles become more interpretive, supervisory, integrative, or boundary-spanning. The strategic issue is whether the organization continues to reproduce the capabilities required for consequential judgment.

Managerial capability follows the same logic. Human managers can remain consequential within algorithmic systems when they translate algorithmic outputs, interpret circumstances that the system does not encode adequately, coordinate exceptions, and intervene outside prescribed routines [18]. Such functions depend on reflexivity and access to relevant

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

information. A nominal human supervisor who no longer understands the work, cannot interrogate the system, or lacks authority to alter its consequences provides little practical protection.

Preservation should therefore be defined functionally. The object is not a historical job description but a capability that the organization may need to exercise under future conditions. Continued human task exposure becomes strategically relevant when it contributes materially to the maintenance of unaided judgment, organizational learning, professional interpretation, or corrective intervention.

The relationship between these mechanisms and future reversibility can be made explicit by separating the capability being preserved from the particular form of human work through which it is maintained.

Table 1 presents the analytical distinctions discussed here.

Table 1. Expertise-preservation mechanisms that determine whether future human re-entry remains credible

Analytical unit	Capability that continued human exposure may preserve	AI use that can remain compatible with preservation	Strategic value for later reallocation	Trade-off if human exposure is retained too broadly	Failure mode if preservation is assumed rather than demonstrated
Unaided decision competence	Ability to perform, calibrate, and correct consequential judgments without relying on an AI recommendation	Retrieval, simulation, drafting, challenge, or bounded advice that does not eliminate recurrent human decision practice	Makes fallback execution more credible if AI becomes unavailable, unreliable, or inappropriate	Foregone speed or quality gains on tasks where AI consistently outperforms unaided judgment	Maintaining nominal human involvement that is too infrequent or passive to sustain competence
Organizational learning richness	Continued exposure to varied actions, outcomes, exceptions, and feedback through which organizational routines are updated	AI-supported analysis where humans still observe and interpret decision consequences	Preserves learning pathways needed to redesign workflows or reassume decision responsibility	Slower standardization and reduced benefits from highly consistent automation	Treating preserved headcount as evidence that organizational learning is still occurring
Professional interpretive capacity	Ability to contextualize ambiguous cases, integrate domain norms, and recognize when formal outputs conflict with professional knowledge	AI use that complements rather than removes opportunities for substantive professional interpretation	Preserves domain-specific judgment that can be mobilized when cases exceed model assumptions	Protecting outdated jurisdictional boundaries or professional routines with declining value	Equating professional status with superior judgment regardless of demonstrated competence

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

Reflexive managerial and expert intervention	Capacity to interrogate outputs, translate them for local conditions, coordinate exceptions, and intervene when consequences become unacceptable	Algorithmic support with genuine access to reasoning, operational context, escalation channels, and intervention rights	Converts formal oversight into actionable control during failure, dispute, or redesign	Additional coordination cost and slower execution where intervention rarely adds value	Symbolic human oversight in which the nominal reviewer lacks knowledge, information, or authority

Commitment Reversibility

Expertise preservation matters strategically because present allocation choices shape the set of credible choices available later. Commitment reversibility refers here to an organization's ability to alter or reverse an AI allocation without first reconstructing the knowledge, routines, authority, and coordination capacity necessary to perform the task differently. This definition extends beyond technical switching costs. A system can be contractually replaceable while the surrounding organization has become practically dependent on it.

The first source of irreversibility is cumulative narrowing. Generative AI can improve the quality of individual creative output while simultaneously increasing similarity across outputs at the collective level [19]. That finding does not establish that AI adoption necessarily reduces organizational strategic diversity, but it illustrates a broader mechanism: local performance improvement can coexist with the narrowing of the experiences or alternatives from which future action is generated. If an organization optimizes repeatedly around the same computational assistance, it may preserve output quality while losing exposure to alternative reasoning pathways.

Learning arrangements can also create irreversible effects before formal dependency is recognized. Research on junior consultants using generative AI shows that digitally capable employees can engage in forms of novice risk work that misidentify or incompletely manage the organizational risks of an emerging technology [20]. Rapid adoption can therefore reshape learning before organizations have developed reliable methods for diagnosing what users must understand about the system itself. A workforce may become proficient at operating AI-supported processes without acquiring the deeper capabilities required to detect when those processes are inappropriate.

The credibility of a fallback arrangement is particularly important. Productive delegation depends on accurate knowledge of relative human and AI capability. Experimental research indicates that humans may delegate poorly when they lack the metaknowledge required to identify when the AI should perform a task and when they should retain it themselves [21]. Preserving authority alone does not solve this problem. An organization can maintain a policy stating that humans may override or replace AI while gradually losing the practiced competence required to know when such intervention is necessary.

Commitment reversibility therefore depends on more than contractual exit rights or technical interoperability. It includes the retention of organizational capacities required to exercise a different allocation. Where those capacities depend on repeated practice, deliberate non-adoption of selected tasks may operate as a form of option preservation. The organization accepts a current opportunity cost in order to keep an alternative operational configuration credible.

That logic has strict limits. Continued human execution is costly when AI delivers sustained benefits and the required fallback capability can be maintained through simulation, rotation, training, external expertise, or other mechanisms that do not require routine non-adoption. Reversibility may also be less important where errors are readily observable, consequences are easily repaired, and alternative suppliers or systems can be substituted without reconstructing

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

organizational knowledge. In those conditions, broad withholding of AI would protect an option whose future value is small.

The strategic concern is therefore not irreversible technology adoption in the abstract. It is capability-coupled commitment: an arrangement becomes difficult to reverse when current deployment reduces the human or organizational capacities needed to enact an alternative later. Deliberate non-adoption is potentially valuable only where the protected capability is genuinely consequential, continued practice materially contributes to its preservation, and cheaper preservation mechanisms are insufficient.

Control over Consequential Tasks

Organizational control over AI-mediated work cannot be inferred from who formally owns the final decision. Algorithmic systems can direct activity, recommend actions, record behavior, evaluate performance, allocate rewards, and progressively replace elements of human judgment [22]. These mechanisms alter practical control because they determine which actor defines the relevant information, frames alternatives, establishes acceptable performance, and initiates intervention. Algorithmic adoption can also reorganize regimes of knowing, gradually shifting which forms of evidence and expertise receive organizational authority [23]. A manager may therefore retain nominal authorization while becoming increasingly dependent on categories, predictions, and workflows whose construction is no longer practically contestable.

The distinction becomes especially important for consequential tasks. Professionals working with opaque AI in diagnostic settings do not engage with outputs uniformly; their willingness and ability to interrogate the system depend on the significance of the judgment, accessible contextual knowledge, and available practices for resolving uncertainty [24]. Institutional conditions can weaken contestability further. In public-sector algorithmic decision making, organizational, social, and legal processes may add layers of blackboxing even when the underlying technology is not literally inaccessible [25]. Effective control therefore requires more than an override button. It requires sufficient information to recognize a reason for intervention, authority to act, and a viable alternative capability through which intervention can be executed.

Transparency can support this arrangement, but it is not an automatic substitute for control. Different transparency mechanisms produce different behavioral consequences; explaining system reasoning and communicating uncertainty need not affect trust or reliance in the same direction [26]. The relevant design problem is consequently conjunctive. Practical control exists when decision makers possess informational access, recognized intervention rights, and maintained execution capability. Removing any one of these elements can convert apparently human-controlled AI into symbolic oversight.

For consequential tasks, deliberate non-adoption becomes plausible where continued AI execution would materially weaken one of these prerequisites and where the resulting loss would be difficult to repair. Where transparency, reliable escalation, and alternative human competence can be maintained under AI use, however, withholding AI has a substantially weaker control rationale.

Strategic Non-Adoption Capability Framework

Prescriptive Design Theory requires organizational prescriptions to specify the context in which an arrangement is intended to operate, the objective it serves, the mechanism through which it should work, the actors who must enact it, and the justificatory knowledge connecting the design to its expected consequences [27]. Deliberate non-adoption therefore cannot be expressed as a general principle that consequential work should remain human. Its object is a narrower capability: the capacity to diagnose when human execution itself preserves strategically relevant resources and to maintain that allocation only while the diagnosed mechanism remains active.

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

The framework contains four linked operations. Task reservation identifies a bounded activity for which continued human execution protects expertise, reversibility, or practical control. Capability maintenance ensures that retained human participation provides enough recurrent exposure to preserve the relevant competence rather than merely satisfying a formal oversight requirement. Controlled coexistence permits AI to support adjacent analytical, retrieval, simulation, or challenge functions where doing so does not remove the protected human capability. Allocation review periodically tests whether technological performance, error recoverability, institutional safeguards, or alternative training mechanisms have changed enough to justify a different arrangement.

These operations place deliberate non-adoption beside, rather than above, competing human–AI configurations. Research on algorithm aversion and appreciation shows that acceptance and resistance vary with task, system, decision structure, and user interpretation [28]. Non-adoption should therefore not be inferred from low trust, preference for human judgment, or resistance to automation. Its distinguishing feature is a documented allocation decision whose preservation mechanism is explicit and whose continuation is conditional.

The capability also has a built-in failure criterion. If an organization cannot identify what continued human execution preserves, cannot demonstrate that practice contributes materially to maintaining it, or can preserve the same capability through a less costly arrangement, the design rationale for non-adoption weakens. Likewise, withholding AI becomes difficult to defend when the arrangement systematically sacrifices substantial decision quality or safety without an offsetting preservation benefit.

The resulting design logic is summarized through the organizational operations that must remain coupled for non-adoption to function as a capability rather than an inertial state.

Table 2 presents the analytical distinctions discussed here.

Table 2. Organizational operations that convert selective AI non-adoption from passive non-use into a reviewable strategic capability

Prescriptive design operation	Article-specific diagnostic question	Mechanism that must remain active	Implementation condition	Principal trade-off	Failure mode requiring redesign
Task reservation	Which bounded consequential activity would lose strategically relevant human capability if routine execution moved to AI?	Continued exposure preserves task-relevant judgment or control capacity	Reserved task is specified narrowly enough to avoid blanket technological exclusion	Foregone AI speed, scale, or consistency	Reservation expands beyond the mechanism that originally justified it
Capability maintenance	Does retained human execution actually sustain expertise needed for future independent action?	Practice, feedback, exception handling, and calibration continue to renew competence	Human participation is substantive and recurrent	Maintaining duplicate or partially redundant capability	Humans remain nominally involved while competence still decays
Controlled coexistence	Which AI functions can be used without displacing the protected capability?	AI complements retained judgment without absorbing the learning or	Functional boundaries between support and substitution are governable	Coordination burden and slower workflow integration	Assistance gradually becomes de facto substitution

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

		authority needed for fallback			without explicit review
Allocation review	Have AI capability, error recoverability, institutional safeguards, or alternative training mechanisms changed the original design problem?	Periodic reassessment preserves the reversibility of the allocation itself	Decision ownership and review triggers remain visible	Review costs and temporary instability	Non-adoption persists after its preservation mechanism has disappeared

Conditions and Propositions

The central argument becomes testable only when deliberate non-adoption can be distinguished from other forms of human retention. Substitutive AI can alter employees' professional role identity particularly where employees lose opportunities to interact with or override the system [29]. That evidence does not demonstrate that preserving human execution will protect professional identity. It does, however, identify interaction and override capacity as consequential features of the allocation itself. A genuine non-adoption capability should therefore produce observable governance arrangements that preserve specified decision practice, rather than merely leaving employees formally accountable after substantive work has migrated elsewhere.

Proposition 1. Deliberate AI non-adoption will have greater strategic value when the expertise required for future independent performance depends strongly on recurrent execution of the reserved task.

The proposition weakens where capability can be maintained through simulation, rotation, external expertise, or other forms of deliberate practice. It also weakens where AI improves human learning rather than displacing it.

Proposition 2. Deliberate AI non-adoption will have greater strategic value when the consequences of erroneous delegation are difficult to detect or reverse and when practical human re-entry would otherwise require rebuilding degraded organizational capability.

This proposition separates consequentiality from simple task importance. A high-value decision whose errors are transparent and readily reversible may provide a weaker case for withholding AI than a less visible process whose cumulative errors are difficult to reconstruct.

Proposition 3. Formal human override rights will provide weaker protection when continued AI dependence erodes the expertise required to recognize when override is warranted.

The discriminating implication is that organizations with similar formal governance may possess very different degrees of practical control because their retained human capabilities differ.

Proposition 4. The case for deliberate non-adoption will weaken as validated AI task fit, error recoverability, effective contestability, and alternative capability-maintenance mechanisms improve.

These propositions prevent the framework from becoming an anti-adoption doctrine. They also create a falsification condition: persistent non-use without an identified preservation mechanism, active governance, and periodic reassessment should be classified as inertia, constraint, resistance, or another form of non-adoption rather than as the strategic capability proposed here.

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

The boundary logic is summarized below by separating nominally similar human authority arrangements according to the conditions that determine whether control remains exercisable.

The relationships discussed here are shown in **Figure 2.**

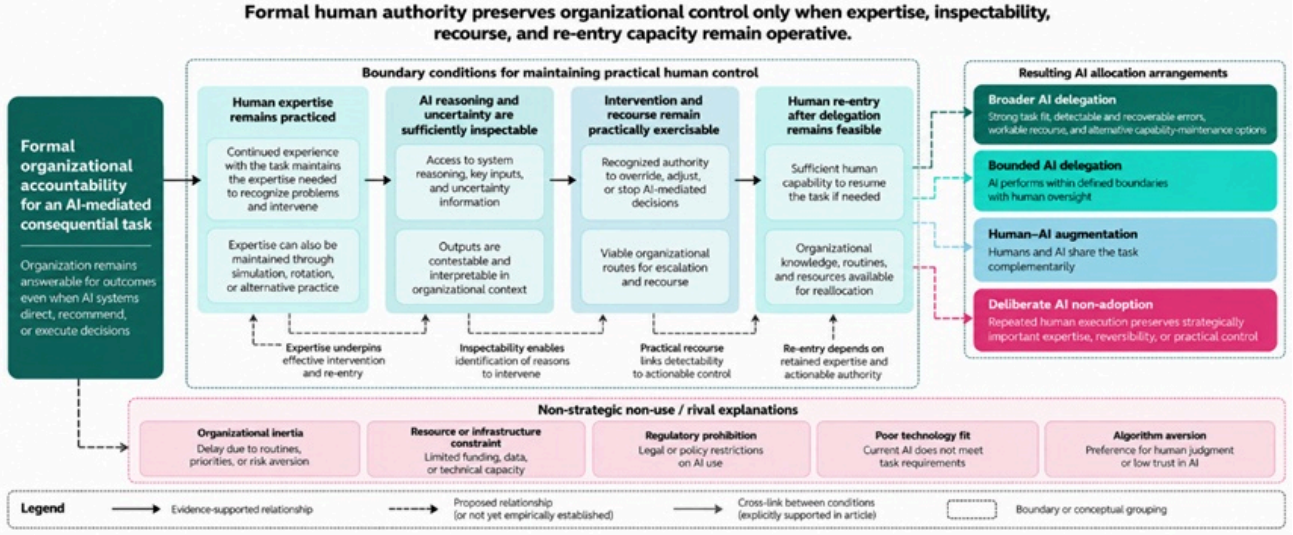

Figure 2. Formal human authority preserves organizational control only when expertise, inspectability, recourse, and re-entry capacity remain operative

Strategy and Governance Implications

The framework shifts AI strategy from maximizing adoption toward maintaining a defensible portfolio of task allocations. This does not imply that organizations should privilege human judgment whenever uncertainty exists. Even strategic evaluation illustrates why configuration matters: individual generative-AI evaluations can be inconsistent or biased while structured aggregation across models can substantially improve their usefulness [30]. Governance should therefore compare arrangements rather than technologies in isolation.

For tasks assigned to AI, organizations should specify who can interrogate the output, who has authority to intervene, and whether those actors retain the competence needed to exercise that authority. For tasks deliberately reserved for human execution, governance should document the capability being preserved and the conditions under which reservation will be reconsidered. This prevents temporary caution from becoming permanent technological avoidance.

Capability maintenance also requires explicit attention to practice. Where independent judgment deteriorates without use, human participation must be frequent and substantive enough to renew competence. Where equivalent capability can be preserved through simulations, rotations, structured exercises, or alternative work, routine non-adoption becomes less defensible. The design principle is preservation of strategic capability, not preservation of existing work arrangements.

Finally, review mechanisms should assess both sides of the trade-off. Organizations should monitor whether withheld AI would now provide substantial validated benefits and whether the protected human capability still contributes to error recognition, organizational learning, re-entry, or control. A strategic non-adoption capability is therefore successful when it enables better reallocation over time—including eventual adoption—not when it maximizes the amount of work kept away from AI.

Conclusion

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

Deliberate AI non-adoption can constitute a strategic capability, but only under restrictive conditions. Its value arises when continued human execution materially preserves expertise needed for consequential judgment, keeps alternative organizational arrangements practically reversible, or sustains the capacity to scrutinize and intervene in AI-mediated decisions. Formal human authority is insufficient when the knowledge required to exercise that authority has disappeared.

The theory therefore rejects both automatic adoption and generalized technological avoidance. Non-adoption is strategically defensible only when it is intentional, bounded, governed, linked to an identifiable preservation mechanism, and periodically reconsidered against alternative arrangements. Where AI performance is reliable, errors are recoverable, contestability is effective, and human capability can be maintained by less costly means, greater AI use is consistent with the same design logic.

The strategic capability is ultimately not abstention from AI. It is the organizational capacity to preserve enough expertise, reversibility, and practical control to decide credibly when AI should be used, when its use should remain bounded, and when human execution must still be kept operational.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 17 Apr 2025 Revised: 22 Jun 2025 Accepted: 07 Jul 2025
Published online: 18 September 2026

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at https://creativecommons.org/licenses/by-nc-sa/4.0/.

References

1. Raisch S, Krakowski S. Artificial intelligence and management: The automation-augmentation paradox. Acad Manage Rev. 2021;46(1):192-210.

https://doi.org/10.5465/amr.2018.0072.

2. Anthony C, Bechky BA, Fayard AL. “Collaborating” with AI: Taking a system view to explore the future of work. Organ Sci. 2023;34(5):1672-94. https://doi.org/10.1287/orsc.2022.1651.

3. Dell’Acqua F, McFowland E III, Mollick E, Lifshitz H, Kellogg KC, Rajendran S, et al. Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. Organ Sci. 2026;37(2):403-23. https://doi.org/10.1287/orsc.2025.21838.

4. Brynjolfsson E, Li D, Raymond L. Generative AI at work. Q J Econ. 2025;140(2):889-942.

5. Papagiannidis E, Mikalef P, Conboy K. Responsible artificial intelligence governance: A review and research framework. J Strateg Inf Syst. 2025;34(2):101885. https://doi.org/10.1016/j.jsis.2024.101885.

6. Grote G, Parker SK, Crowston K. Taming artificial intelligence: A theory of control-accountability alignment among AI developers and users. Acad Manage Rev. 2026;51(2):278-99. https://doi.org/10.5465/amr.2023.0117.

7. Shrestha YR, Ben-Menahem SM, von Krogh G. Organizational decision-making structures in the age of artificial intelligence. Calif Manage Rev. 2019;61(4):66-83. https://doi.org/10.1177/0008125619862257.

8. Konsynski BR, Kathuria A, Karhade PP. Cognitive reapportionment and the art of letting go: A theoretical framework for the allocation of decision rights. J Manag Inf Syst. 2024;41(2):328-40. https://doi.org/10.1080/07421222.2024.2340830.

9. Vanneste BS, Puranam P. Artificial intelligence, trust, and perceptions of agency. Acad Manage Rev. 2025;50(4):726-44. https://doi.org/10.5465/amr.2022.0041.

10. Cameron LD. The making of the “good bad” job: How algorithmic management manufactures consent through constant and confined choices. Adm Sci Q. 2024;69(2):458-514. https://doi.org/10.1177/00018392241236163.

11. Townsend DM, Hunt RA, Rady J, Manocha P, Jin JH. Are the futures computable? Knightian uncertainty and artificial intelligence. Acad Manage Rev. 2025;50(2):415-40. https://doi.org/10.5465/amr.2022.0237.

12. Grønsund TR, Aanestad M. Augmenting the algorithm: Emerging human-in-the-loop work configurations. J Strateg Inf Syst. 2020;29(2):101614. https://doi.org/10.1016/j.jsis.2020.101614.

13. Cruz I. Expert-AI pairings: Expert interventions in AI-powered decisions. Inf Organ. 2024;34(4):100527. https://doi.org/10.1016/j.infoandorg.2024.100527.

14. Bauer K, Nofer M, Henrich B, Drachsler H, Hinz O. The dependency dilemma: How machine learning decision aids can undermine skill growth. Bus Inf Syst Eng. 2026;68(5):1121-46. https://doi.org/10.1007/s12599-026-01014-z.

15. Sarala RM, Post C, Doh J, Muzio D. Advancing research on the future of work in the age of artificial intelligence (AI). J Manag Stud. 2025;62(5):1863-84. https://doi.org/10.1111/joms.13195.

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

16. Balasubramanian N, Ye Y, Xu M. Substituting human decision-making with machine learning: Implications for organizational learning. Acad Manage Rev. 2022;47(3):448-65. https://doi.org/10.5465/amr.2019.0470.

17. Scarbrough H, Chen Y, Patriotta G. The AI of the beholder: Intra-professional sensemaking of an epistemic technology. J Manag Stud. 2025;62(5):1885-913. https://doi.org/10.1111/joms.13065.

18. Leavitt K, Barnes CM, Shapiro DL. The role of human managers within algorithmic performance management systems: A process model of employee trust in managers through reflexivity. Acad Manage Rev. 2025;50(4):745-67. https://doi.org/10.5465/amr.2022.0058.

19. Doshi AR, Hauser OP. Generative AI enhances individual creativity but reduces the collective diversity of novel content. Sci Adv. 2024;10(28):eadn5290. https://doi.org/10.1126/sciadv.adn5290.

20. Kellogg KC, Lifshitz H, Randazzo S, Mollick E, Dell'Acqua F, McFowland E III, et al. Novice risk work: How juniors coaching seniors on emerging technologies such as generative AI can lead to learning failures. Inf Organ. 2025;35(1):100559. https://doi.org/10.1016/j.infoandorg.2025.100559.

21. Fügener A, Grahl J, Gupta A, Ketter W. Cognitive challenges in human-artificial intelligence collaboration: Investigating the path toward productive delegation. Inf Syst Res. 2022;33(2):678-96. https://doi.org/10.1287/isre.2021.1079.

22. Kellogg KC, Valentine MA, Christin A. Algorithms at work: The new contested terrain of control. Acad Manage Ann. 2020;14(1):366-410. https://doi.org/10.5465/annals.2018.0174.

23. Pachidi S, Berends H, Faraj S, Huysman M. Make way for the algorithms: Symbolic actions and change in a regime of knowing. Organ Sci. 2021;32(1):18-41. https://doi.org/10.1287/orsc.2020.1377.

24. Lebovitz S, Lifshitz-Assaf H, Levina N. To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using AI for medical diagnosis. Organ Sci. 2022;33(1):126-48. https://doi.org/10.1287/orsc.2021.1549.

25. Kronblad C. When justice is blind to algorithms: Multilayered blackboxing of algorithmic decision-making in the public sector. MIS Q. 2024;48(4):1637-62.

26. Vössing M, Kühl N, Lind M, Satzger G. Designing transparency for effective human-AI collaboration. Inf Syst Front. 2022;24(3):877-95. https://doi.org/10.1007/s10796-022-10284-3.

27. Gregor S, Chandra Kruse L, Seidel S. Research perspectives: The anatomy of a design principle. J Assoc Inf Syst. 2020;21(6):1622-52. https://doi.org/10.17705/1jais.00649.

28. Jussupow E, Benbasat I, Heinzl A. An integrative perspective on algorithm aversion and appreciation in decision-making. MIS Q. 2024;48(4):1575-90.

29. Strich F, Mayer AS, Fiedler M. What do I do in a world of artificial intelligence? Investigating the impact of substitutive decision-making AI systems on employees' professional role identity. J Assoc Inf Syst. 2021;22(2):304-24. https://doi.org/10.17705/1jais.00663.

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies

30. Doshi AR, Bell JJ, Mirzayev E, Vanneste BS. Generative artificial intelligence and evaluating strategic decisions. Strateg Manage J. 2025;46(3):583-610. https://doi.org/10.1002/smj.3677.

© 2026 The Author(s). | Published by Institute for Management, Business, and Accounting Studies