ORIGINAL RESEARCH **Open access**

The Expiry of Corporate AI Assurance When Systems Change Faster Than Boards Can Renew Evidence of Their Reliability and Organizational Fitness

Michael Brenner[1*], Thomas Berger[2], Julia Weber[3], Elena Rossi[4]

Abstract

Corporate AI assurance is commonly treated as evidence that a system is sufficiently reliable, governed, and organizationally fit at the time of assessment. That position becomes unstable when the model, its data, its deployment context, or the processes surrounding its use change on different temporal cycles from those through which assurance evidence is generated and renewed. This article develops a Temporal / Multiple-Clocks Theory of corporate AI assurance. It conceptualizes assurance as condition-dependent evidence whose organizational validity persists while circumstances material to the supported claim remain sufficiently stable. The analysis distinguishes technological, data, contextual, process, monitoring, escalation, oversight, and institutional clocks and examines how their synchronization or desynchronization alters feasible governance responses. Its central contribution is the concept of assurance expiry as a relational state arising when material validity-threatening change outpaces the organizational processes capable of detecting, evidencing, escalating, and authorizing reassessment. Expiry is therefore neither a universal chronological deadline nor an inevitable consequence of every system modification. It depends on change materiality, observability, reversibility, delegation, and evidence-production latency. This account reframes board oversight around the currency of the evidentiary basis for organizational reliance on changing AI systems.

Keywords Model drift, Temporal governance, Corporate AI assurance, Organizational AI governance, Responsible artificial intelligence, Board oversight

*Correspondence:
Michael Brenner
michael.brenner@univie.ac.at
[1] Department of Information Systems, University of Vienna, Vienna, Austria
[2] Department of Management, University of Bern, Bern, Switzerland
[3] Department of Information Systems, Free University of Berlin, Berlin, Germany
[4] Department of Management, University of Bologna, Bologna, Italy

Introduction

Artificial intelligence governance increasingly concerns how responsibilities, controls, review practices, and escalation mechanisms are organized across the AI lifecycle [1]. AI simultaneously reshapes managerial work through changing combinations of automation and augmentation, thereby altering where decision authority and responsibility for evaluating system behavior reside [2]. Together, these developments create a temporal governance problem: the object being governed may change while the evidence supporting organizational reliance on it remains anchored to an earlier system state.

Organization theory provides a useful basis for examining this problem because organizational activities do not unfold under one common clock. Work and organizing exhibit different rhythms, deadlines, cycles, episodes, and horizons

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

whose interaction can create coordination problems even when each activity is internally coherent [3]. Temporal work further indicates that organizational actors can shape temporal assumptions and arrangements through practices that connect past experience, present action, and anticipated futures [4]. Corporate AI assurance can consequently be analyzed as a relationship among interacting clocks.

The issue is not simply that AI changes quickly. Some deployed systems remain technically stable, whereas others are frequently modified. Data may change without model retraining; a system may be introduced into a new business context without changing its code; workflow automation may expand while measured model performance remains stable. Conversely, frequent technical modification may remain governable when monitoring, evidence production, and authorized reassessment occur at a compatible pace.

This article therefore asks when the evidence supporting corporate reliance on AI remains current and when temporal mismatch makes that evidence insufficiently representative of the system actually being governed. Temporal / Multiple-Clocks Theory is used to distinguish system-change processes from evidence-production, escalation, and oversight processes. The resulting contribution is conditional: elapsed time alone does not invalidate assurance, and every technical change does not require board intervention. The theoretical problem concerns the conditions under which system state and assurance state cease to correspond.

Corporate AI Assurance as Time-Bounded Evidence

AI assurance derives meaning from the claim, system boundary, risk, and evidentiary object being assessed. Broader approaches to assurance are important because a narrowly technical assessment can omit harms or dependencies arising elsewhere in development, deployment, and organizational use [5]. Scope therefore has a temporal consequence. Evidence remains informative only while the conditions material to the claim it supports continue to characterize the system being used.

Organizational AI governance provides the institutional arrangements through which such judgments are allocated and acted upon [6]. Ethics-based auditing supplies a more specific evidentiary mechanism: an audit evaluates defined claims against selected criteria and evidence, while its conclusions remain constrained by its scope and evaluative frame [7]. An assurance conclusion can accordingly be interpreted as evidence supported under specified conditions, even when organizational reporting compresses those conditions into a simpler status designation.

Operational auditing also depends on organizational routines. Industry evidence shows that abstract AI principles must be translated into practices involving coordination, communication, assessment, and organizational change [8]. Accountability further requires identifiable actors, standards, answerability, and authority to respond [9]. Evidence may therefore become organizationally stale before anyone formally declares it invalid if material changes occur without reaching an actor capable of initiating reassessment.

The object of assurance must also be bounded carefully. Technical reliability, regulatory acceptability, and organizational fitness are related but distinct. Analysis of trustworthy AI under the European Union AI Act illustrates the conceptual difference between trustworthiness and the regulatory acceptability of risk [10]. A stable compliance designation consequently cannot establish by itself that the original evidence continues to answer the organization’s present reliability and fitness questions.

This article defines an assurance validity interval as the contingent period during which existing assurance evidence remains materially applicable to the AI system as currently used. The interval has no predetermined duration. Its continuation depends on correspondence between the conditions represented in the evidence and the conditions that now govern the system. Chronological age may increase without expiry when relevant conditions remain stable; a material change can shorten the interval abruptly when it affects what the original assessment established.

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

Sources of AI System Change

The governance environment around AI is itself changeable. Mapping of AI-governance initiatives shows heterogeneous arrangements distributed across institutions and instruments [11]. Comparative research on public organizations likewise indicates that governance implementation can depend on different configurations of organizational and institutional conditions [12]. Standards, responsibilities, and escalation expectations may therefore change around a system even when its technical configuration remains constant.

Technology supplies a second clock. Generative AI demonstrates how model capabilities and available uses can evolve as models and deployment arrangements develop [13]. Yet organizational AI capability is broader than the underlying model, incorporating complementary resources and organizational competencies [14]. Changes in integration, expertise, infrastructure, or authorized use may alter the practical system covered by assurance without being reducible to a new model version.

Data constitute another moving layer. Data-governance arrangements influence the conditions under which trustworthy AI is developed and operated [15]. These arrangements include decision rights, accountabilities, policies, standards, and control mechanisms [16]. Changes in data sources, provenance, accessibility, composition, or governance may therefore alter the evidentiary environment of an otherwise unchanged model.

Four principal sources of change follow: model/capability change, data change, deployment-context change, and process-integration change. Institutional change can modify the criterion against which all four are judged. Their clocks need not coincide. The distinction between these sources is important because each can disrupt a different assumption embedded in existing assurance.

The temporal differences are consolidated in **Table 1.**

Table 1. Distinct change clocks that can alter the object of corporate AI assurance

Change clock	Principal change	Assurance-relevant shift	Observable organizational signal	Temporal significance
Model/capability	Model version, enabled functions, autonomy, configuration	Operative capability moves beyond the examined capability state	New release, configuration change, expanded tool access	Changes the technical object represented by existing evidence
Data	Composition, provenance, availability, feature relationships	Current inputs differ materially from assessed inputs	Source substitution, distribution change, altered missingness or population	Changes the informational conditions supporting assurance
Deployment context	Market, geography, population, task, jurisdiction	The same system is used under materially different conditions	New business unit, user group, jurisdiction, or decision setting	Changes where and for whom the assurance claim is applied
Process integration	Human review, workflow sequence, override rights, downstream action	Organizational consequences of outputs change	Workflow redesign, automation expansion, changed human review	Changes how model behavior becomes organizational action
Institutional criterion	Regulation, internal policy, assurance	The question evidence must answer changes	New rule, policy revision, assurance requirement,	Changes the governing criterion

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

	standard, risk expectation		changed risk appetite	even if the system is unchanged

These clocks can interact without forming a single sequence. A context change may expose a data mismatch; process redesign may amplify the consequences of an unchanged prediction; a new institutional requirement may make previously sufficient evidence incomplete. Assurance is therefore conditioned by temporal coupling between the changing system and the processes that establish what remains known about it.

When material change is detected and evidence is refreshed at a compatible pace, correspondence between assurance and deployment can be maintained. If relevant conditions change before monitoring and reassessment catch up, an evidence-applicability gap can emerge without any requirement that the system has catastrophically failed. **Figure 1** represents this architecture.

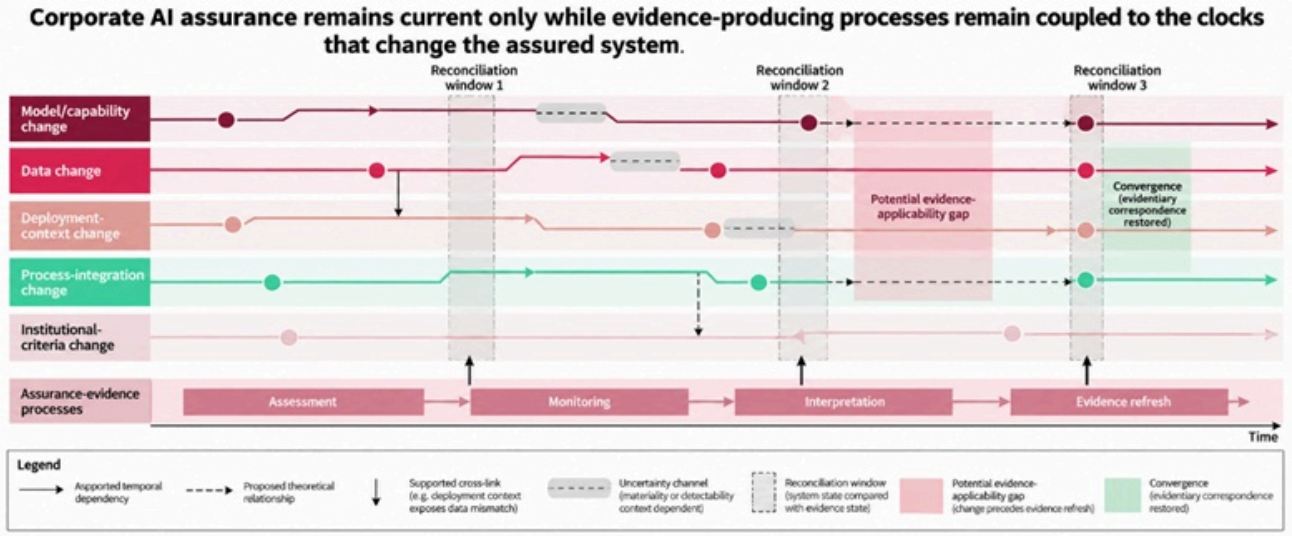

Figure 1. Corporate AI assurance remains current only while evidence-producing processes remain coupled to the clocks that change the assured system

Model, Data, Context, and Process Drift

The preceding section identifies where change originates. Drift concerns how divergence from an earlier assessed state becomes observable. Concept-drift research establishes that relationships represented in data streams can change over time, challenging stationary assumptions embedded in learned models [17]. Such drift may also be local, affecting particular features, regions, or populations instead of the full operating environment [18]. Assurance can consequently become less representative for one use while remaining informative for another.

Detection has its own clock. Reviews of drift detection and adaptation show that methods respond differently to different drift structures and assumptions [19]. Feature-level approaches further indicate that superficially similar changes can arise from distinct underlying relationships [20]. A governance statement that “the model has drifted” therefore contains insufficient information for assurance renewal. What changed, where it changed, and whether the affected relationship is material to the supported assurance claim are separate questions.

Monitoring techniques can detect evolving information properties, but detection is not simultaneous with change. Entropy-based methods provide one approach for identifying distributional change in streams [21]. Postmarket research on medical AI likewise shows that distribution-shift surveillance can identify relevant changes while remaining sensitive to the

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

type and detectability of the shift [22]. The system can therefore occupy a changed state before the organization has adequate evidence to characterize it.

Context creates a different form of divergence. Evidence on model transportability demonstrates that performance can differ across sites when local data-generating conditions vary [23]. Operational work on translational machine learning also emphasizes that workflow integration, data connections, monitoring, and maintenance continue after initial deployment [24]. Although these findings arise from healthcare settings, they establish a bounded mechanism relevant to corporate theory: assurance may depend on model, context, and process jointly.

Adaptive deployment makes the distinction still more important. Dynamic AI systems may require continuing evaluation because aspects of the operative system can change after deployment [25]. Longitudinal evidence from deployed clinical models shows that data and performance shifts can occur over time and differ across use cases and sites [26]. These studies do not establish a universal corporate expiry period; they demonstrate why stable validity cannot simply be presumed after initial assessment.

For the present theory, **drift** denotes a material divergence between conditions represented in existing assurance evidence and conditions governing current system behavior or organizational use. Model, data, context, and process drift remain separate analytical categories because each changes a different part of the evidence-to-deployment relationship. Their governance consequence depends on materiality, observability, reversibility, and whether the divergence falls within existing delegated authority.

Board-Level Assurance Renewal

A temporal theory of assurance must specify who can respond when evidence may no longer represent the contemporary system. Integrative AI-governance research shows that risks differ and therefore invite differentiated governance responses [27]. Corporate assurance renewal should accordingly be organized around the significance of change and the authority needed to respond, rather than one review cadence applied uniformly to every AI system.

Increasing AI autonomy complicates this allocation because autonomous capability can alter established relationships among action, agency, and accountability in corporate governance [28]. Organizational decision theory also identifies multiple ways of distributing decision authority between humans and AI [29]. Board-level assurance renewal therefore does not imply that directors should inspect each model update. The relevant governance question concerns which changes may be handled within delegated operational authority and which alter the evidentiary basis on which board-level reliance is justified.

Organizational case evidence demonstrates that continuing AI governance can be institutionalized through lifecycle roles, committees, review routines, and oversight structures [30]. This makes separation between fast operational clocks and slower oversight clocks possible without eliminating accountability. Specialist functions may monitor continuously, produce new evidence, and resolve predefined classes of change. Escalation occurs when materiality, uncertainty, exposure, or authority requirements exceed those delegated boundaries.

The design problem is two-sided. Slow escalation can permit organizational reliance to continue after evidence has ceased to represent relevant current conditions. Excessive escalation can congest oversight and redirect board attention toward operational changes that specialist governance arrangements are equipped to manage. Temporal fitness therefore depends on coupling decision rights to change materiality while maintaining a traceable path through which evidence can move from operational monitoring to formal assurance renewal.

This allocation is summarized in **Table 2.**

Table 2. Governance clocks connecting operational evidence refresh with board-level assurance renewal

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

Governance clock	Principal actor	Evidence or decision object	Normal authority	Escalation condition	Characteristic temporal failure
Operational monitoring	Model owners, data/ML operations, risk teams	Drift signals, incidents, workflow changes	Detect, investigate, contain	Signal may affect an assured condition	Material change remains unnoticed or buried in monitoring noise
Evidence production	Validation, audit, compliance, domain specialists	Re-testing, contextual evidence, control effectiveness	Refresh or qualify the evidentiary record	Prior evidence no longer represents the current state adequately	Evidence arrives after organizational reliance has already changed
Management escalation	Responsible executive, AI governance committee, risk leadership	Materiality, exposure, reversibility, uncertainty	Resolve within delegated authority or escalate	Change exceeds delegation or risk boundary	Material issues wait below the oversight level or minor issues are over-escalated
Board/oversight renewal	Board or designated committee	Consolidated assurance status and unresolved material change	Renew, qualify, suspend, or redirect organizational reliance	Prior board-level basis for reliance is materially altered	Oversight either rests on stale assumptions or becomes congested with operational detail
Institutional/regulatory	Regulator, standard setter, court, external assurance body	Duties, standards, external findings	Alter the governing evidentiary criterion	External requirement or accepted assurance basis changes	Internal evidence no longer answers the governing assurance question

The multiple-clocks perspective therefore places the board within an assurance system whose components legitimately operate at different speeds. The central governance difficulty arises when those differences cease to be complementary and become temporal misalignment: system conditions move, evidence follows later, escalation follows later still, and organizational reliance continues in the interval. The section on assurance expiry and evidence staleness examines when that interval constitutes assurance expiry rather than ordinary and tolerable governance latency.

Assurance Expiry and Evidence Staleness

Temporal misalignment becomes theoretically consequential when organizational processes remain synchronized with conditions that have already changed. Research on managerial temporal orientations shows that actors operating within the same organization can work from different interpretations of relevant pasts and futures [31]. Temporal miscoupling similarly describes situations in which an established practice remains attached to a temporal environment whose conditions are declining or changing [32]. Applied to AI assurance, this produces a specific problem: organizational reliance can continue under an evidentiary state that still represents an earlier deployment condition.

Trustworthiness further spans technical, organizational, and institutional levels [33]. A reliability signal at one level therefore cannot automatically preserve assurance across the others. Organizational turning points may also alter how

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

prior evidence and anticipated futures inform present action [34]. A material AI-system change can function as such a turning point when it alters a condition upon which an existing assurance claim depends.

Longitudinal post-deployment evidence demonstrates that prior validation can become insufficient as operating conditions evolve [35]. Postmarket governance research likewise identifies challenges created by changing AI/ML systems after authorization [36]. Aggregate assessment can additionally conceal localized underperformance affecting specific subgroups [37]. These findings support a non-binary conception of evidence staleness.

An assurance-staleness interval begins when a material validity-threatening change has occurred but adequate refreshed evidence has not yet been produced and authorized. Assurance expiry occurs when continued reliance during that interval can no longer be justified by the evidentiary basis previously accepted. The transition is conditional: stable conditions may preserve older evidence, while a consequential change can invalidate part of an assurance claim rapidly.

Temporal AI Assurance Framework

The temporal framework treats corporate AI assurance as a relation among eight distinguishable clocks: model, data, deployment context, process integration, monitoring/evidence production, management escalation, board or oversight renewal, and institutional criteria. Multiple temporalities can coexist without failure when their differences remain functionally coordinated [3]. Temporal work allows organizations to reshape some of those relations through routines, escalation rules, review windows, and delegated authority [4].

The decisive variable is therefore relative timing. Assurance remains current when evidence-production and decision processes respond before material divergence undermines the claims on which organizational reliance depends. Audit scope determines which changes matter to a particular claim [7]. Operational governance determines whether emerging evidence reaches an actor authorized to respond [8].

Table 3 distinguishes the principal temporal configurations generated by this logic.

Table 3. Temporal configurations of corporate AI assurance under asynchronous change and renewal

Temporal configuration	Relationship among clocks	Evidentiary state	Governance response
Coupled stability	Relevant system conditions remain stable relative to monitoring and review	Existing evidence remains applicable	Continue monitoring within existing authority
Detected adaptive change	Change precedes review but is rapidly observed and evidenced	Assurance temporarily qualified but recoverable	Refresh evidence within delegated governance
Hidden divergence	Material change precedes effective detection	Evidence appears current while representativeness has weakened	Intensify monitoring and investigate applicability
Escalation lag	Change is known, but evidence or authority moves more slowly	Staleness interval persists	Escalate to the level holding renewal authority
Renewed correspondence	Reassessment catches up with the changed system state	Evidentiary basis again matches present conditions	Renew, narrow, condition, or suspend reliance

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

Renewal Triggers and Governance Propositions

The framework implies that renewal should be triggered principally by material change in an assurance-relevant condition, rather than calendar age alone. Drift research demonstrates why change may occur at different locations and with different observability [17]. Post-deployment surveillance further shows that detection itself is conditional on the character of the shift [22].

The first proposition is therefore that assurance-renewal urgency increases when a detected change affects a condition material to an existing assurance claim. The second is that equivalent technical changes can produce different assurance outcomes when evidence-generation, escalation, or authorization latency differs. Human–AI decision structures make such variation plausible because organizations allocate decision rights differently [29].

A third proposition concerns observability and reversibility. Low-observability changes increase the risk that reliance persists through an undetected staleness interval, whereas highly observable and reversible changes can more readily remain within delegated control. A fourth proposition limits board involvement: more frequent board review is not inherently superior. Governance frameworks support differentiated responses to differentiated risks [27], and lifecycle governance can distribute responsibilities across specialized organizational actors [30].

These propositions generate the boundary-condition map in **Figure 2.**

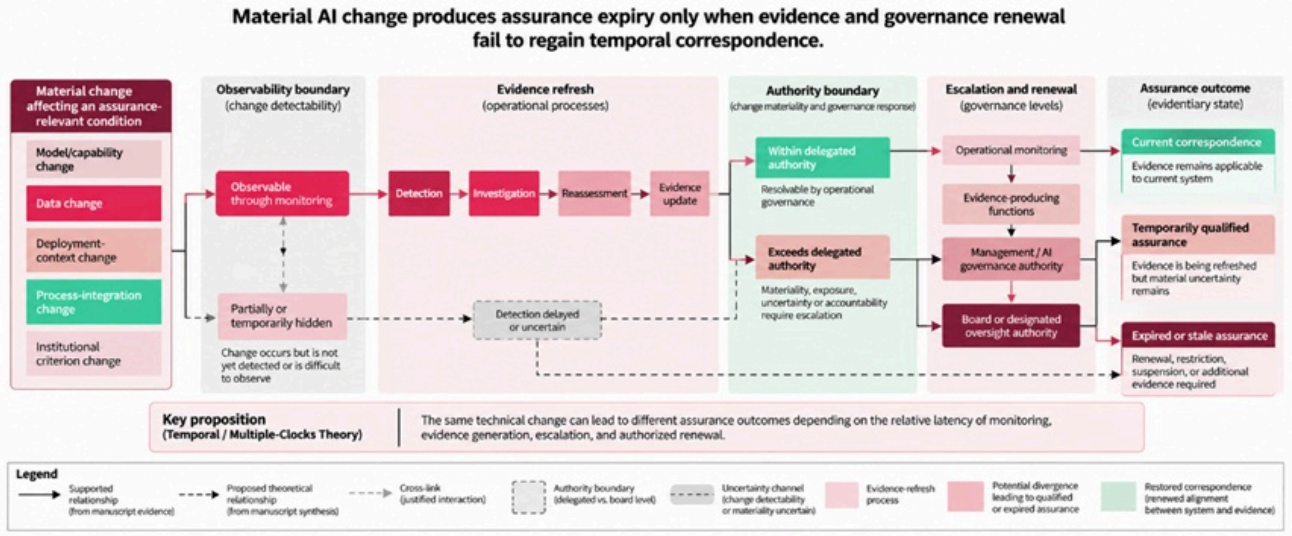

Figure 2. Material AI change produces assurance expiry only when evidence and governance renewal fail to regain temporal correspondence

Corporate Governance Implications

Boards need visibility into the currency and scope of assurance evidence, rather than only a point-in-time performance statement. Responsible-AI governance already places AI within organizational structures and lifecycle practices [1]. Assurance practice adds the requirement to communicate what has been assessed and where its limits lie [5].

This does not require boards to become operational monitors. Accountability can remain distributed while authority and escalation responsibilities are explicit [9]. Delegated renewal is therefore compatible with board oversight when materiality boundaries, reporting duties, and authority to restrict or suspend reliance are defined.

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

Corporate AI assurance should consequently be designed as an evidence-renewal system. Monitoring functions detect change; technical and domain specialists determine whether it matters to an existing claim; management resolves changes within delegated authority; and the board or designated committee receives matters that alter the basis on which organizational reliance was authorized. Governance quality then depends partly on whether those clocks remain sufficiently coupled to prevent material evidentiary divergence from persisting unnoticed.

Conclusion

Corporate AI assurance does not expire merely because time passes. It becomes stale when material changes in the system, data, deployment context, organizational process, or governing criterion advance beyond the evidence and authority supporting continued reliance.

Temporal / Multiple-Clocks Theory makes this distinction visible by separating change, detection, evidence production, escalation, and renewal. These processes can legitimately operate at different speeds. The critical condition is whether their relative timing preserves correspondence between the system being used and the evidence on which governance relies.

The article therefore reframes assurance as condition-dependent and renewable. Where material change is observable and rapidly absorbed within legitimate delegated authority, assurance can remain current without repeated board intervention. Where change outruns detection, evidence production, or authorized reassessment, an assurance-staleness interval emerges. Board-level renewal becomes necessary when that mismatch alters the evidentiary basis on which corporate reliance was justified.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 03 Jun 2025 Revised: 08 Aug 2025 Accepted: 24 Aug 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

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

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. 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.

2. 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.

3. Omidvar O, Hadjimichael D, Burke GT, Pyrko I, Chia R. Temporal patterns in management: Integrating perspectives on rhythms of work and organizing. Acad Manage Ann. 2025;19(2):861-902. https://doi.org/10.5465/annals.2024.0081.

4. Bansal P, Reinecke J, Suddaby R, Langley A. Temporal work: The strategic organization of time. Strateg Organ. 2022;20(1):6-19. https://doi.org/10.1177/14761270221081332.

5. Thomas C, Roberts H, Mökander J, Tsamados A, Taddeo M, Floridi L. The case for a broader approach to AI assurance: Addressing “hidden” harms in the development of artificial intelligence. AI Soc. 2025;40(3):1469-84. https://doi.org/10.1007/s00146-024-01950-y.

6. Mäntymäki M, Minkkinen M, Birkstedt T, Viljanen M. Defining organizational AI governance. AI Ethics. 2022;2(4):603-9. https://doi.org/10.1007/s43681-022-00143-x.

7. Mökander J, Morley J, Taddeo M, Floridi L. Ethics-based auditing of automated decision-making systems: Nature, scope, and limitations. Sci Eng Ethics. 2021;27(4):44. https://doi.org/10.1007/s11948-021-00319-4.

8. Mökander J, Floridi L. Operationalising AI governance through ethics-based auditing: An industry case study. AI Ethics. 2023;3(2):451-68. https://doi.org/10.1007/s43681-022-00171-7.

9. Novelli C, Taddeo M, Floridi L. Accountability in artificial intelligence: What it is and how it works. AI Soc. 2024;39(4):1871-82. https://doi.org/10.1007/s00146-023-01635-y.

10. Laux J, Wachter S, Mittelstadt B. Trustworthy artificial intelligence and the European Union AI Act: On the conflation of trustworthiness and acceptability of risk. Regul Gov. 2024;18(1):3-32. https://doi.org/10.1111/rego.12512.

11. Attard-Frost B, Brandusescu A, Lyons K. The governance of artificial intelligence in Canada: Findings and opportunities from a review of 84 AI governance initiatives. Gov Inf Q. 2024;41(2):101929. https://doi.org/10.1016/j.giq.2024.101929.

12. de Almeida PGR, dos Santos Júnior CD. Artificial intelligence governance: Understanding how public organizations implement it. Gov Inf Q. 2025;42(1):102003. https://doi.org/10.1016/j.giq.2024.102003.

13. Feuerriegel S, Hartmann J, Janiesch C, Zschech P. Generative AI. Bus Inf Syst Eng. 2024;66(1):111-26. https://doi.org/10.1007/s12599-023-00834-7.

14. Mikalef P, Gupta M. Artificial intelligence capability: Conceptualization, measurement calibration, and empirical study on its impact on organizational creativity and firm performance. Inf Manag. 2021;58(3):103434. https://doi.org/10.1016/j.im.2021.103434.

15. Janssen M, Brous P, Estevez E, Barbosa LS, Janowski T. Data governance: Organizing data for trustworthy artificial intelligence. Gov Inf Q. 2020;37(3):101493. https://doi.org/10.1016/j.giq.2020.101493.

16. Abraham R, Schneider J, vom Brocke J. Data governance: A conceptual framework, structured review, and research agenda. Int J Inf Manage. 2019;49:424-38. https://doi.org/10.1016/j.ijinfomgt.2019.07.008.

17. Lu J, Liu A, Dong F, Gu F, Gama J, Zhang G. Learning under concept drift: A review. IEEE Trans Knowl Data Eng. 2019;31(12):2346-63. https://doi.org/10.1109/TKDE.2018.2876857.

18. Aguiar GJ, Cano A. A comprehensive analysis of concept drift locality in data streams. Knowl Based Syst. 2024;289:111535. https://doi.org/10.1016/j.knosys.2024.111535.

19. Arora S, Rani R, Saxena N. A systematic review on detection and adaptation of concept drift in streaming data using machine learning techniques. Wiley Interdiscip Rev Data Min Knowl Discov. 2024;14(4):e1536. https://doi.org/10.1002/widm.1536.

20. Yang L, Cheng J, Luo Y, Zhou T, Zhang X. Detecting and rationalizing concept drift: A feature-level approach for understanding cause-effect relationships in dynamic environments. Expert Syst Appl. 2025;260:125365. https://doi.org/10.1016/j.eswa.2024.125365.

21. Sun Y, Mi J, Jin C. Entropy-based concept drift detection in information systems. Knowl Based Syst. 2024;290:111596. https://doi.org/10.1016/j.knosys.2024.111596.

22. Koch LM, Baumgartner CF, Berens P. Distribution shift detection for the postmarket surveillance of medical AI algorithms: A retrospective simulation study. NPJ Digit Med. 2024;7(1):120. https://doi.org/10.1038/s41746-024-01085-w.

23. Lasko TA, Strobl EV, Stead WW. Why do probabilistic clinical models fail to transport between sites. NPJ Digit Med. 2024;7(1):53. https://doi.org/10.1038/s41746-024-01037-4.

24. Poddar M, Marwaha JS, Yuan W, Romero-Brufau S, Brat GA. An operational guide to translational clinical machine learning in academic medical centers. NPJ Digit Med. 2024;7(1):129. https://doi.org/10.1038/s41746-024-01094-9.

25. Rosenthal JT, Beecy A, Sabuncu MR. Rethinking clinical trials for medical AI with dynamic deployments of adaptive systems. NPJ Digit Med. 2025;8(1):252. https://doi.org/10.1038/s41746-025-01674-3.

26. Cabanillas Silva P, Sun H, Rezk M, Roccaro-Waldmeyer DM, Fliegenschmidt J, Hulde N, et al. Longitudinal model shifts of machine learning-based clinical risk prediction models: Evaluation study of multiple use cases across different hospitals. J Med Internet Res. 2024;26:e51409. https://doi.org/10.2196/51409.

27. Wirtz BW, Weyerer JC, Kehl I. Governance of artificial intelligence: A risk and guideline-based integrative framework. Gov Inf Q. 2022;39(4):101685. https://doi.org/10.1016/j.giq.2022.101685.

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

28. Ahdadou M. From augmentation to autonomy: Artificial intelligence and the destabilization of agency in corporate governance. Futures. 2026;182:103873. https://doi.org/10.1016/j.futures.2026.103873.

29. 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.

30. Kim JY, Hasan A, Kueper J, Tang T, Hayes C, Fine B, et al. Establishing organizational AI governance in healthcare: A case study in Canada. NPJ Digit Med. 2025;8(1):522. https://doi.org/10.1038/s41746-025-01909-3.

31. Sasaki I, Kotosaka M, De Massis A. When top managers' temporal orientations collide: Middle managers and the strategic use of the past. Organ Stud. 2024;45(6):825-53. https://doi.org/10.1177/01708406241236604.

32. Ortiz Casillas S, Huising R. Temporal miscoupling: The challenges and consequences of enacting a practice in decline. Organ Sci. 2025;36(1):288-312. https://doi.org/10.1287/orsc.2021.16026.

33. Lahusen C, Maggetti M, Slavkovik M. Trust, trustworthiness and AI governance. Sci Rep. 2024;14(1):20752. https://doi.org/10.1038/s41598-024-71761-0.

34. Hernes T, Feuls M. Organizing in the folding of time: Shaping organizational change trajectories at turning points. Organ Stud. 2024;45(9):1325-48. https://doi.org/10.1177/01708406241272935.

35. Kopanitsa G. Validation is not enough: Longitudinal evidence of post-deployment fragility in clinical AI systems. PLOS Digit Health. 2026;5(7):e0001534. https://doi.org/10.1371/journal.pdig.0001534.

36. Babic B, Cohen IG, Stern AD, Li Y, Ouellet M. A general framework for governing marketed AI/ML medical devices. NPJ Digit Med. 2025;8(1):328. https://doi.org/10.1038/s41746-025-01717-9.

37. Subbaswamy A, Sahiner B, Petrick N, Pai V, Adams R, Diamond MC, et al. A data-driven framework for identifying patient subgroups on which an AI/machine learning model may underperform. NPJ Digit Med. 2024;7(1):334. https://doi.org/10.1038/s41746-024-01275-6.

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