

Research on the Collaborative Governance Framework for Organizational Information Resources in the Era of Generative Large Language Models: An Integrative Design of Actors, Instruments, and Mechanisms

Chengjun Xie*

Liaoning Communication University, Shenyang 111000, China

*Corresponding email: 271246@qq.com

Abstract

Generative large language models continue to redefine how organizations produce, verify, and utilize information resources. Existing governance arrangements assume that data assets are relatively stable and that system behavior is observable. This assumption no longer holds for outputs that are generated probabilistically, cannot be traced to a source, and are produced jointly by humans and machines. Therefore, this paper proposes an integrative framework in which actors, instruments, and mechanisms interlock. A generative model occupies multiple positions simultaneously within the governance relationship: It is subject to governance arrangements, it is employed as an instrument of governance, and it is a force capable of shaping decisions while remaining beyond accountability. Authority and responsibility are first redistributed among information professionals, business units, technology vendors, and algorithmic systems. Instruments are considered next, and three distinct types must not be conflated: technical devices that discipline inputs, means of governing via generative models, and arrangements that oversee the first two. At the level of mechanisms, structural, procedural, relational, and consequential designs are integrated. The primary challenge lies not in the capability of these systems, but in how long-standing weaknesses - namely ambiguous ownership, evidence that escapes review, and judgment displaced by procedure - are amplified to the point where they can no longer be ignored.

Keywords

Generative large language models, Accountability, Collaborative governance, Governance arrangements

Introduction

Generative large language models have transitioned from experimental apparatuses in the laboratory to infrastructure that organizations rely on for daily operations. They increasingly mediate the relationship between an organization and its own stock of knowledge [1]. What first leaves existing governance arrangements flat-footed is generativity itself. Outputs are synthesized rather than retrieved and cannot be traced back along a lineage graph to a source. The combination of opacity and fluency makes such outputs appear entirely reasonable while providing no basis for contesting them. More awkward still, a generative model simultaneously occupies three positions: It is an object to be governed, an instrument available for use, and a force that can sway decisions without being accountable for them, as

illustrated in Figure 1. This paper, therefore, asks what constitutes the governing body, which instruments are actually available, and whether the existing mechanisms can withstand scrutiny [2].

Three points define the contribution. First, the three positions a generative model occupies are treated as a single condition rather than as separate problems, which explains why arrangements addressing only one position tend to fail. Second, actors, instruments, and mechanisms are integrated into a unified design rather than listed as parallel pillars, allowing a change in one element to be evaluated in relation to the other two. Finally, the couplings between them are treated as the fundamental unit of design, based on the view that most failures are failures of connection. The argument proceeds as follows:

The research foundations and the identified gap are first established. The framework is then decomposed into its constituent elements, reassembled into three layers, and

examined through the couplings between them. Subsequently, its application and boundary conditions are discussed.




Figure 1. Three positions of generative large language models in the governance relation and the central tensions.

Research foundations

From data governance to artificial intelligence governance

Data governance is not a one-off act of control. It is a structured allocation of decision rights, and those rights must be assigned deliberately [3]. Artificial intelligence (AI) governance has since been defined at the organizational level, distinguishing it from corporate governance, information technology (IT) governance, and data governance [4]. The field has been reorganized around the questions of who governs, what is governed, and how governance proceeds [5]. Meanwhile, principles for trustworthy artificial intelligence have been extended from individual organizations to multi-organizational networks.

This extension carries two shifts. The first is the object of governance. Data governance addresses an asset whose properties can be described before use, whereas generative output is produced at the moment of use and cannot be fully specified in advance. The second is the locus of control. Much of what determines an output is

settled outside the organization that uses it - in the selection of the model, the version in force, and the corpus on which retrieval depends. Consequently, an arrangement confined to internal procedures reaches only part of the process [6].

Instruments, overreliance, and the research gap

Confronted with systems they cannot fully explain, organizations respond through sociotechnical envelopment: Delimiting interaction boundaries, selecting training corpora with care, and controlling the sources of inputs and outputs [7]. Managers tend to place greater trust in advice generated by artificial intelligence, leading to overreliance [8]. According to one influential view, accountability should attach to the concrete act of using a generative system [9]. Each of these studies advances the discussion, yet none treats actors, instruments, and mechanisms as mutually constitutive elements.

What remains missing is a framework for holding these three elements together. Studies of instruments seldom ask who is authorized to deploy them, and studies of actors seldom ask what a decision right is worth when the evidence required to exercise it was never captured. The

consequence is a familiar asymmetry: Governance arrangements accumulate where documentation is easy but thin out where judgment is required. This paper, therefore, takes the coupling between the elements - rather than any single element in isolation - as the unit of design.

Conceptual deconstruction: Actors, instruments, and mechanisms

Actors

Information professionals are accountable for data integrity and retention periods. Business units wield a different kind of power: They decide which content may enter the system and for what purposes an output may be used. Technology vendors control model selection and version updates, exercising a de facto governing authority that is seldom documented. Generative systems bear no responsibility themselves, yet their outputs exert causal force, a force that must be represented on the governance map.

Two complications follow. The first is drift. A role that held decision rights at the time of system deployment may retain its title after the actual work has shifted elsewhere; consequently, the map of authority describes a practice that no longer exists. The second is asymmetry of knowledge. Those best positioned to judge whether an output is sound are often not those authorized to release it. Therefore, the actor layer must specify how that knowledge informs the decision, not merely who makes it.

Instruments

The first type of instrument disciplines inputs. Corpus governance, access control, and retrieval-augmented architectures belong to this category. The second governs through the use of generative models, as seen in automated classification and the detection of policy deviations. The third type supervises the first two, utilizing provenance capture, review workflows, and audit trails. These three types fail in distinct ways. Input-side instruments degrade silently. Generative instruments can dampen scrutiny by making unsupported assertions. Supervisory instruments fail when their outputs are taken as guarantees rather than as evidence. A fourth type is sometimes proposed: instruments that act on the market for models through contractual terms and procurement conditions. These are real, but they operate

at a distance from the information resource itself. This paper sets them aside so that the three primary kinds remain distinct. More importantly, these kinds are not interchangeable. Strengthening a review process does not rectify a corpus that should never have been admitted, and tightening access controls does not validate an unsupported assertion.

Mechanisms

Structural mechanisms allocate decision rights and establish bodies with clear mandates. Procedural mechanisms specify how resources enter a system and how outputs are verified. Relational mechanisms sustain the shared understanding necessary for coordination across professional disciplines. Consequential mechanisms set performance levels and feed information about deviations back to the other three types. Investment tends to concentrate on the first two, as they are the easiest to document.

The neglect of relational mechanisms carries an easily overlooked cost. Coordination across professional groups is what enables a reviewer to recognize a problem that no rule anticipated. This coordination is sustained through opportunities for joint judgment, not through documentation alone. The consequential type, for its part, is frequently reduced to efficiency indicators, which report throughput while remaining silent on whether the outputs produced are ones the organization would defend.

An integrative framework

The actor layer defines who holds decision rights and who is accountable. The instrument layer provides the means required for governance. The mechanism layer is responsible for coordinating actors, disciplining instruments, and connecting governance to outcomes. These three layers are not successive stages. Instruments presuppose authorized actors, and mechanisms are what enable an allocation of authority to operate, see Figure 2. Sequencing is what distinguishes this design from a checklist. Because instruments presuppose authorized actors, an organization that acquires supervisory instruments before determining who may act upon them acquires documentation rather than governance. Because mechanisms enable an allocation of authority to operate, a decision-rights matrix that no procedure consults remains inert. The layers are better understood as conditions for one another than as components assembled in sequence.

Collaboration carries three meanings within this design. The most basic occurs between people. Authority is distributed across professional groups, and the arrangement requires joint judgment rather than deference to an owner. Information resources and the models that utilize them increasingly exist outside any single hierarchy, so cross-organizational collaboration cannot be avoided. Human-machine collaboration is the most easily misinterpreted. It represents limited delegation, not partnership. The system is entrusted only with what it performs reliably;

whether a consequence is reversible and whether it is acceptable remain judgments made by people [10]. The framework is deliberately sparing in its prescriptions. It identifies the positions that must be filled and the couplings that must be maintained, leaving their fulfillment to the organization. This restraint follows from the subject matter rather than preference, since the points of contact between generative systems and information resources vary enough across organizations for a detailed prescription to be misleading.




Figure 2. A three-layer integrative framework for collaborative governance: actors, instruments, and mechanisms.

The actor layer

Accountability cannot be transferred. This is the first requirement of the actor layer. A system can bring about consequences yet cannot be held to account; therefore, wherever a generative system contributes materially to an information product, a named human role must retain responsibility. De facto power must also be made explicit. The choice of model configuration, prompt libraries, and retrieval corpora is itself a governing decision. The human-machine relationship should be configured for

augmentation rather than substitution, and the boundary of independent machine action is determined by whether the consequences can be reversed.

One further requirement concerns vendors. Their power over model selection and versioning is exercised through decisions that appear technical and are rarely routed through any governance body. Bringing that power into the open does not require the vendor to join the organization's arrangements. It requires the organization to record which configuration was in force when a given

output was produced and who approved the change.

The instrument layer

The deployment order of the three kinds cannot be inverted. Input-side instruments must come first because, once unvetted corpus material has entered a system, no downstream arrangement can recover the loss. Generative instruments may be enabled only where output passes established review, and supervisory instruments are a precondition for enabling them, not an optional accessory. Overreliance is the default tendency, so review procedures should require justification rather than invite confirmation. Explainability is still unavailable, and assurance can therefore be sought only at three boundaries: What enters the system, what leaves it, and who approved it.

Capacity, not availability, is usually what constrains the supervisory kind. Provenance capture and review workflows exist in most organizations that have attempted them. What is missing is reviewer time sufficient to operate them at the volume of output actually produced. A realistic deployment therefore estimates review load against that volume and scales the generative kind to the review capacity rather than the reverse.

The mechanism layer

Structural mechanisms yield a decision-rights matrix - long advocated in data governance research - together with a body empowered to suspend deployment. Procedural mechanisms require outputs to be verified before they enter records that carry consequences and to leave a trace of whether a generative system was used. Relational mechanisms create opportunities for joint judgment across professional boundaries. Consequential mechanisms set the expected standard. Their function is correction rather than reporting upward.

A body empowered to suspend deployment is the structural mechanism most often omitted, and its absence is felt precisely when it is needed. Without it, the procedural requirement to verify an output carries no sanction, and a reviewer who finds a problem has no recourse other than escalation to someone who may lack the standing to act.

Couplings between layers

An instrument is effective only to the extent that authority stands behind it. Logs that no one can act upon produce documentation, not governance. Instruments

produce evidence, and mechanisms decide whether that evidence is examined and by whom. Mechanisms also decide how investment is allocated: If consequential indicators capture efficiency alone, the relational work that sustains judgment is marginalized. The three layers therefore must be designed together.

These couplings also explain why failures resist localization. An output that causes harm may be traceable to a corpus decision made months earlier, to a review that lacked the capacity to examine it, and to a set of indicators that rewarded the volume of production. Each layer performed as designed; what failed was the connection between them.

Discussion

The process begins with mapping. Identify the points at which generative systems interact with information resources and assign decision rights at each contact point; this typically reveals the extent of informally exercised power. The subsequent step involves assessing whether supervisory capacity matches the scale of deployment, with the shortfall usually lying on the supervision side. The sequence concludes with mechanism design. Verification procedures and escalation paths must be established, along with a set of diagnostic rather than summative indicators.

Two caveats arise from applying the framework in this order. Mapping contact points tends to uncover informal power that is not documented, and there is often an impulse to formalize it immediately; doing so before supervisory capacity is assessed transforms an informal arrangement into an unenforced one. Estimating supervision, in turn, produces uncomfortable figures because the review workload scales with output rather than with the number of systems deployed.

Evaluation deserves separate consideration because the chosen indicators determine the nature of the arrangement. If only efficiency is recorded, review is perceived as a cost and is gradually reduced; if the proportion of outputs examined and the proportion reversed upon examination are also recorded, supervision gains its own standing. A small set of diagnostic indicators, reported to the body authorized to suspend deployment, is more useful than a comprehensive dashboard that no one is obligated to act upon.

One objection should be addressed directly. It may be

argued that the framework adds a layer of procedure to organizations already lacking capacity and that generative systems were adopted precisely to reduce such burdens. The response is that the burden is not created by the framework but revealed by it. Review of generated output already occurs informally, typically when someone decides whether to trust it; the framework proposes that this judgment be located, recorded, and assigned to someone with the authority to act on it.

This framework is not a universal prescription. It presupposes that decision-making roles within the organization can be identified by name. Flat organizations often lack this condition, and organizations with limited resources may find the costs outweigh the benefits.

Two further conditions deserve mention. The framework assumes that an output can be attributed to a decision. When generative systems are chained, with one model's output serving as another's input, attribution becomes difficult and the actor layer loses its anchor. It also assumes that records are retained long enough for review to occur afterward. If retention periods were set for stable assets, the evidence of generative use may expire before anyone examines it.

The implications for professional practice are concrete. Information professionals are already asked whether a generated output can be relied upon, and the framework suggests that this question be answered based on three recorded facts: What entered the system, who approved the release, and whether the output was examined afterward. If those three cannot be produced, the arrangement has not been implemented, regardless of what its documents state.

Cost is the objection most frequently raised at this point, and it deserves a straightforward answer. Supervision does scale with output, so an organization that increases its review proportion without limiting where generative systems are applied will find the requirement unsustainable. The framework accommodates this by making the scope of application a governance decision rather than a technical one: Where review capacity is limited, the range of uses should be narrowed before the standard of review is lowered.

Conclusion

Generative models blur the line between what is <https://www.wonford.com/>

governed and what governs, thereby introducing into the governance process an element that affects outcomes without bearing accountability. The tripartite division proposed in earlier work is extended here to the domain of information resources, specifying the couplings among actors, instruments, and mechanisms. At the practical level, regulation should focus on how decision rights are allocated within an organization and on the retention of evidence. The former determines whether a responsible party can be identified once an output causes harm; the latter determines whether the reasoning can be reconstructed afterward.

The validity of this framework has not been tested against primary data, which is its main limitation. Case studies tracking deployment as it unfolds could transform coupling strength into measurable indicators.

Several lines of inquiry follow. One is empirical and would examine whether the couplings described here hold in organizations that have already deployed generative systems, using deployment records rather than stated policy. Another is comparative and would investigate whether arrangements designed for regulated domains are transferable to organizations where the same outputs circulate without external scrutiny. A third concerns the instruments themselves and would explore the extent to which provenance can be captured when the model in use is replaced by a vendor without notice.

Funding

This work was not supported by any funds.

Acknowledgements

The author would like to show sincere thanks to those technicians who have contributed to this research.

Conflicts of Interest

The author declares no conflict of interest.

References

- [1] Taeiagh, A. (2025) Governance of generative AI. *Policy and Society*, 44(1), 1-22.
- [2] Wang, X., Liu, B., Bao, L. (2026) Generative artificial intelligence policies: information governance. *Management Decision*, 64(9), 4074-4110.
- [3] Abraham, R., Schneider, J., Vom Brocke, J. (2019) Data governance: a conceptual framework, structured review, and research agenda.

- International Journal of Information Management*, 49, 424-438.
- [4] Mäntymäki, M., Minkkinen, M., Birkstedt, T., Viljanen, M. (2022) Defining organizational AI governance. *AI and Ethics*, 2(4), 603-609.
- [5] Schneider, J., Abraham, R., Meske, C., Vom Brocke, J. (2023) Artificial intelligence governance for businesses. *Information Systems Management*, 40(3), 229-249.
- [6] Janssen, M., Brous, P., Estevez, E., Barbosa, L. S., Janowski, T. (2020) Data governance: Organizing data for trustworthy artificial intelligence. *Government Information Quarterly*, 37(3), 101493.
- [7] Asatiani, A., Malo, P., Rådberg Nagbøl, P., 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.
- [8] Keding, C., Meissner, P. (2021) Managerial overreliance on AI-augmented decision-making processes: How the use of AI-based advisory systems shapes choice behavior in R&D investment decisions. *Technological Forecasting and Social Change*, 171, 120970.
- [9] Susarla, A., Gopal, R., Thatcher, J. B., Sarker, S. (2023) The Janus effect of generative AI: Charting the path for responsible conduct of scholarly activities in information systems. *Information Systems Research*, 34(2), 399-408.
- [10] Chowdhury, S., Budhwar, P., Wood, G. (2024) Generative artificial intelligence in business: towards a strategic human resource management framework. *British Journal of Management*, 35(4), 1680-1691.