

THE GOVERNABILITY GAP: STRUCTURAL LIMITS OF GOVERNANCE UNDER CONDITIONS OF SYSTEM COMPLEXITY

Tatiana Arden
University of Ljubljana
ardenform@gmail.com

Abstract

Governance systems are failing not simply because they are poorly designed, but because they are operating beyond their capacity. Across domains such as climate, pandemics, financial stability, and AI oversight, institutional responses increasingly fail to match the complexity of the systems they seek to govern. Explanations centered on political will, resources, or administrative inefficiency describe symptoms, but not the underlying structure.

This paper argues that contemporary governance failure is systemic. It arises from a limit condition defined by a mismatch between system-generated variety and governing capacity. Building on the law of requisite variety, the viable system model, and insights from cognitive science, the paper introduces the concept of the governability gap: a structural imbalance between the variety produced by complex socio-technical systems and the cognitive and organizational variety available to institutions.

The paper proposes a core mechanism: institutions do not simply compensate for bounded rationality; they often scale it. As system complexity increases, governance architectures can amplify cognitive bias, short time horizons, and information bottlenecks. This creates a recursive dynamic in which institutional interventions increase system variety while also degrading the capacity to absorb it.

This dynamic is conceptualized as a viability threshold. Beyond a certain level of complexity, governance systems can no longer maintain requisite variety, coordinate across subsystems, or sustain adaptive coherence. This threshold is emergent rather than fixed. It arises from interactions between human cognitive constraints, institutional design, and system-level feedback loops. While not fully predictable, it is diagnosable.

The contribution is threefold: it reframes governance failure as a predictable outcome of variety imbalance; develops a diagnostic framework for identifying approaching or exceeded viability thresholds; and outlines a design agenda for institutions operating under complexity constraints, including modularity, distributed sense-making, and adaptive variety alignment.

Keywords

governability gap, requisite variety, viable systems, Substitution-Cost Indicator (SCI), boundary-constructability

1 | Introduction

Some governance problems do not only resist intervention. They change in response to it.

This is visible in climate governance, pandemic response, financial regulation, digital platform control, and AI oversight. More data enters the system. More rules are written. More expert attention is applied. The immediate problem may be partly addressed, but the field then reorganizes. New evasions appear. New risks are displaced elsewhere. More monitoring becomes necessary. The next round of governance starts from a more complex position than the previous one.

The easy explanation is that institutions lack political will, resources, competence, or coordination. These factors matter. They do not explain the repeated structure of the failure. A different explanation is needed for cases where institutional effort grows and yet governability declines.

This paper treats that pattern as a systems problem. More precisely, it treats it as a problem of variety, modelling, and viability. Ashby's Law of Requisite Variety states that only variety can absorb variety

The Governability Gap

(Ashby, 1956). A regulator must command enough variety to distinguish and respond to the states of the system it seeks to regulate. Beer's viable system model extends this into organizational cybernetics: a viable system must maintain enough internal variety, coordination, adaptation, and identity to survive in its environment (Beer, 1972, 1979, 1985). Conant and Ashby add the missing hinge: every good regulator must contain a model of the system it regulates (Conant & Ashby, 1970).

That hinge matters. Governance never acts on the governed system in its full complexity. It acts through a model. It uses categories, procedures, indicators, roles, targets, eligibility rules, risk classifications, and reporting formats. The model is necessary. A system that tried to represent everything would not be able to act. But the model also selects. It decides what can be seen, what can be counted, what can be rewarded, and what remains outside the regulatory field.

The governability gap begins there. It is not only the distance between environmental complexity and institutional capacity. It is also the distance between the living system and the model through which governance tries to act on it.

The paper develops one diagnostic mechanism inside that broader gap. It argues that governance models tend to see visible output more clearly than they see the conditions that make output possible. Two different losses follow. First, the model undercounts maintenance: repair, recovery, care, prevention of breakdown, and the reproduction of capacity. Second, it tends to proxy foundations: trust, meaning, legitimacy, and professional judgment. These foundations reduce the cost of coordination, but direct measurement often changes them. A trust score is not trust. A loyalty metric is not legitimacy. A procedure for initiative is not judgment.

The main operational proposal is the Substitution-Cost Indicator, or SCI. SCI does not try to measure foundations directly. It tracks the cost of their substitutes: more control, more audit, more incentives, more enforcement, more procedural specification, more symbolic production of cohesion. If a system keeps output stable only by paying more and more for coordination, it may be approaching a viability threshold. The paper then connects this signal to boundary-constructability. Foundations can be protected only where there is a boundary within which defection can be seen, addressed, and limited. Yet a boundary can work in two ways. It can become self-sustaining, reducing the need for external control over time. Or it can remain externally sustained, requiring more control to hold the same behavior in place. The difference is not ethical rhetoric. It is feedback.

This is the diagnostic question the paper adds to the governability gap: when institutions add organization, does the cost of coordination fall, or does it rise?

2 | Method and Scope

This is a conceptual paper. It does not report an empirical study. Its method is a systems-theoretical synthesis with a diagnostic aim.

The argument draws on five bodies of work. The first is cybernetics: Ashby's requisite variety, Conant and Ashby's good regulator theorem, and Beer's viable system model (Ashby, 1956; Beer, 1972, 1979, 1985; Conant & Ashby, 1970). The second is bounded rationality and institutional cognition, especially Simon's account of limited human and organizational decision capacity (Simon, 1955, 1957). The third is work on legibility, measurement, and proxy failure, especially Scott, Goodhart, and Strathern (Goodhart, 1975; Scott, 1998; Strathern, 1997). The fourth is second-order cybernetics and autopoiesis, where the observer is not outside the observed system (Maturana & Varela, 1980; von Foerster, 1981). The fifth is commons governance and sociotechnical design, especially Ostrom's design principles and Cherns's principle of minimal critical specification (Cherns, 1976; Ostrom, 1990).

The paper does not claim to invent the problem of governability. That field is already well established, especially in work on interactive governance, which distinguishes the governing system from the system-to-be-governed and asks under what conditions governance can occur at all (Kooiman, 2003). Nor does the paper claim that measurement is always harmful. That would be too easy, and wrong. Some invisible work needs better recognition. Some systems need stronger control. Some metrics prevent real harm.

The claim is narrower. It says that governance needs to distinguish between two kinds of invisibility. One can often be partly corrected by better recognition. The other is damaged by direct conversion into

The Governability Gap

indicators. This distinction then makes a diagnostic question possible: if we cannot measure a foundation directly, can we detect its depletion through the rising cost of substitutes?

The paper uses organizational examples because they are easier to see than whole societies. The point, however, is not limited to firms, schools, or universities. Organizations are a middle scale. They make visible what is also present in public institutions: models, feedback loops, control costs, boundaries, and the conditions under which adaptive capacity is preserved or drawn down.

3 | The Governability Gap as a Variety Problem

The starting point is variety.

A governance system must distinguish enough states of the governed system to respond adequately. If the governed system can take more meaningful states than the regulator can recognize and answer, control becomes partial. Sometimes that is acceptable. No governance system needs total control. The problem begins when the regulator cannot distinguish the states that matter for viability.

This is already a hard limit. Human cognition is bounded. Institutional cognition is bounded too, even when it appears larger than individual cognition. Institutions do not automatically overcome bounded rationality by adding people, procedures, and data. Often they reproduce it in another form. They filter. They simplify. They shorten time horizons. They build information bottlenecks. They create reporting hierarchies where weak signals are lost or translated into categories that no longer carry the original meaning.

This is what I mean by saying that institutions scale bounded rationality. They do not merely compensate for it. In complex settings, they may amplify it.

The good regulator theorem makes the mechanism sharper. The regulator must contain a model of the regulated system (Conant & Ashby, 1970). For social governance, this means that institutions act through models of persons, groups, firms, citizens, patients, students, workers, claimants, users, risks, and needs. The living subject is too various to govern directly. So governance builds a legible subject and acts on that. Scott's work on legibility shows why this is especially strong at scale (Scott, 1998). Large systems prefer categories that travel. They need information that can be recorded, compared, aggregated, audited, and acted upon from a distance. Local knowledge is harder to move. Tacit judgment is harder to count. Maintenance is harder to display than output. Foundations are harder still.

The governability gap, then, is not simply too much complexity. It is a patterned mismatch between system-generated variety and the model through which governance tries to absorb it. The model simplifies in order to act. The simplification changes the field. Actors adapt to the model. Some capacities are undercounted. Some foundations are proxied. The regulator sees the failure and adds more specification. That new specification becomes the next environment to which actors adapt.

The gap regenerates.

This regenerative pattern matters for viability. A system remains viable only while it can absorb environmental and internal variety without destroying the conditions of its own adaptation. Once governance interventions create more variety than the system can absorb, a threshold may be approached. The threshold is not a fixed number. It is a condition. It appears when the system can no longer maintain requisite variety, coordinate its subsystems, or sustain adaptive coherence.

Such a threshold cannot be predicted with precision. It can, however, be diagnosed indirectly.

The rest of the paper develops one route for doing so.

4 | Models, Maintenance, and Foundations

Governance models are selective. They usually see output first.

Output is visible because it leaves a positive trace. A report is filed. A patient is processed. A publication appears. A ticket is closed. A budget is spent. A target is met. These units can be counted, compared, and used in decisions.

The conditions that make output possible are less visible. They often do not appear as new units. They appear as continued capacity. The system still works. People still cooperate. The machine still runs. The group does not break down. The professional still knows when not to follow the script.

This paper separates two kinds of invisibility within that field.

4.1 | Type One Invisibility: Maintenance

Maintenance is activity that preserves the capacity of a system to continue. It includes repair, care, recovery, prevention of failure, reproduction of skill, retention of institutional memory, and protection of working conditions.

This use of maintenance also draws on work on repair, infrastructure, and invisible work (Graham & Thrift, 2007; Jackson, 2014; Star & Strauss, 1999).

Maintenance often takes the form of a non-event. A failure does not happen. A conflict does not escalate. A person does not burn out. A machine does not break. A team does not lose coordination.

A kindergarten teacher is a clear example. Much of the work consists in preventing the group from crossing a threshold: the child does not collapse, the conflict does not become a crisis, the group does not fall into chaos. The best result may look like nothing happened. Standard reporting systems often miss this, because they register activities, plans, complaints, readiness, incidents, or measurable learning outcomes. They do not easily register prevented breakdown.

A technical analogy is preventive maintenance. A reserve system, a stress test, or a replaced component may also produce value as a non-event. The outage did not occur. The system did not fail. In technical settings the non-event is often easier to connect to a failure model. In social systems it is harder, because the work is distributed across attention, timing, professional judgment, relationships, and small adjustments. Still, maintenance can often be made partly visible without destroying it. Not perfectly. Enough to protect it. A school can attend to group stability, staff exhaustion, transition quality, and prevented incidents. A hospital can monitor handover quality, recovery time, coordination load, and staff fatigue. A university can recognize mentoring, peer review, program repair, and the informal work that keeps research and teaching possible.

This is Type One invisibility. The object is undercounted. The response is to widen the model, carefully.

4.2 | Type Two Invisibility: Foundations

Type Two invisibility concerns foundations.

A foundation is a system property that lowers the cost of coordination but is damaged when treated as a direct target. The main examples here are trust, meaning, legitimacy, and professional judgment.

The definitions are functional.

Trust is the capacity to coordinate without constant verification. Meaning is the capacity to sustain action without continuous external incentive. Legitimacy is the capacity of a rule to produce compliance without a proportional rise in enforcement. Professional judgment is the capacity to make context-sensitive decisions without full prior specification.

These are not ornamental values. They are control-cost reducers.

When trust is present, a team does not verify every small action. When meaning is present, effort does not need to be bought moment by moment. When legitimacy is present, rules do not need constant coercive backing. When professional judgment is present, the organization does not need to pre-script every case.

The problem is that foundations cannot be made visible in the same way as maintenance. A trust score is not trust. A loyalty metric is not legitimacy. An engagement index is not meaning. A checklist of good judgment is not judgment.

Direct measurement creates a proxy. At that point, the familiar problem of indicator corruption begins: indicators used for control or decision-making become subject to pressure and may distort the processes they were meant to monitor (Campbell, 1976; Goodhart, 1975). In audit settings, this is not only a technical failure of measurement; it can reshape the institution around the indicator itself (Strathern, 1997). In foundation cases, the damage is deeper. The measure may replace the object. People learn to perform trust, loyalty, engagement, or judgment in the form required by the metric.

The proxy grows. The foundation may decline.

Second-order cybernetics helps explain why. Observation is not outside the system. It is part of the system's operation (von Foerster, 1981). In social systems, the observed can read the observation, adapt to it, and change the meaning of the observed behaviour. Maturana and Varela's work on autopoiesis also warns

The Governability Gap

against treating living systems as if their organization were fully available to external control (Maturana & Varela, 1980).

Type Two invisibility therefore calls for restraint. The question is not how to measure the foundation directly. The question is how to protect it, and how to detect its depletion without replacing it with a proxy. That is where SCI enters.

5 | The Substitution-Cost Indicator

If a foundation cannot be measured directly, it may still leave traces when it weakens.

The trace is cost.

A working foundation makes coordination cheaper. When it weakens, the system must purchase substitutes. Trust declines, and verification expands. Meaning declines, and incentives multiply. Legitimacy declines, and enforcement rises. Professional judgment declines, and procedural specification grows. Cohesion declines, and symbolic production of unity becomes louder.

The Substitution-Cost Indicator, or SCI, tracks this movement.

SCI does not measure foundations directly. It tracks the cost of substitutes for depleted foundations relative to the system's own declared core output.

The Substitution-Cost Indicator can be expressed as Equation (1).

$$SCI_t = \frac{\sum_{i=1}^n w_i C_i(t)}{Q_{core}(t)} \quad (1)$$

In Equation (1), $C_i(t)$ denotes observed substitution costs at time t , including control, audit, compliance, incentives, procedural specification, verification, enforcement, retention programs, and symbolic cohesion work. The weights w_i are empirical. $Q_{core}(t)$ refers to the system's declared core output.

This is an operational heuristic, not a derived mathematical model. Its value lies in tracking change over time. SCI compares a system with itself. It does not rank different systems as if they had the same function. The definition of Q_{core} matters. It must be set relative to the declared function of the system itself. In a hospital, clinical care is core output; extra reporting time per case may be a substitution cost. In an audit firm, verification is part of core output; internal verification added to compensate for loss of trust may still be a substitution cost. In a university, teaching and research are core outputs; expanding administrative work per unit of teaching or research may indicate a rising substitution cost.

The point is not that all control is bad. That would be useless. Some control is necessary. Some audit prevents harm. Some enforcement protects the vulnerable. SCI asks a narrower question: is the system increasingly paying for coordination that a foundation once provided at lower cost?

Consider a university department. It can count publications, teaching hours, grants, student evaluations, and formal procedure completion. It sees those well. It has more difficulty seeing trust among colleagues, informal mentoring, willingness to share time, professional reputation, and the collective judgment that keeps academic work meaningful. If these foundations weaken, the department may respond with more reporting, more mandatory meetings, more formal mentoring frameworks, more quality assurance procedures, more performance reviews, more culture programs.

Output may remain stable. The cost of producing it rises.

That rise is the signal.

The same pattern appears in organizations that maintain delivery metrics by adding layers of approval, in public agencies that maintain compliance by adding enforcement, and in teams that maintain performance by increasing incentives while informal commitment fades. SCI does not prove that the foundation is depleted. It marks where investigation should begin.

This is the diagnostic value. It gives governance a way to see the indirect cost of what its own model cannot directly see.

6 | Boundary-Constructability

SCI can show that a foundation is being replaced by substitutes. It does not yet show whether the foundation can be protected.

For that, we need boundaries.

Foundations are vulnerable because foundation-preserving behavior often loses in the short term. A person who maintains trust, professional standards, or shared resources carries costs now. Others benefit later. An actor who extracts from the same foundation can win immediately.

This is a commons problem. Olson showed the collective action difficulty: as groups grow, individual incentives often drift away from collective maintenance (Olson, 1965). Ostrom showed that commons do not survive through goodwill alone. They depend on institutional arrangements: clear boundaries, monitoring, graduated sanctions, conflict-resolution mechanisms, local rule-making, and nested governance (Ostrom, 1990).

The same logic applies to organizational foundations.

A team cannot protect trust if opportunistic behavior is invisible. A profession cannot protect judgment if the organization rewards only speed and compliance. A public institution cannot protect legitimacy if those who draw down legitimacy face no consequence. A university cannot protect intellectual quality if every incentive favors countable throughput.

A foundation needs a boundary.

Boundary-constructability is the capacity to build and maintain a frame within which foundation-damaging behavior can be identified, addressed, and limited without destroying cooperation itself.

This is not simply the presence of rules. Rules can be distant, rigid, and poorly matched to the field. Boundary-constructability includes who is inside the regime of obligation, what counts as defection, whether defection can be observed, whether response is possible, whether sanctions are legitimate, and whether participants can adjust rules as conditions change.

Boundary-constructability is high when membership is clear, scale is workable, behavior is observable, sanctions are proportionate, and local knowledge remains active. It falls when the system is too large, too open, too anonymous, or too dependent on distant proxies.

This explains why trust-based governance often works in one setting and fails in another. Trust did not suddenly become naive. The boundary that made trust workable may not have scaled.

Cherns's principle of minimal critical specification is relevant here (Cherns, 1976). It does not say no rules. It says specify what is essential and leave the rest to local autonomy. This is a way to preserve variety inside a boundary. It works only when the boundary can hold. Otherwise autonomy becomes risk transfer. The center withdraws specification without creating the conditions under which responsible self-regulation can survive.

Boundary-constructability is therefore a scope condition for adaptive governance.

6.1 | Boundary Feedback: Self-Sustaining and Externally-Sustained Boundaries

Boundary-constructability answers one question: can a boundary be built?

A second question is just as important: what happens after it is built?

A boundary can protect a foundation. It can also replace it. The difference is visible in feedback over time. A boundary is self-sustaining when it reduces the need for external control. After the boundary is introduced, participants increasingly reproduce the desired behaviour from within the system. The rule becomes part of the system's own regulation. Coordination becomes cheaper. External monitoring can fall. A boundary is externally sustained when the same behaviour requires rising control, checks, sanctions, or enforcement. The visible form of order remains, but only because the external loop grows. Coordination becomes more expensive. In SCI terms, the boundary is not rebuilding the foundation. It has become one more substitute for it.

This distinction avoids a false opposition between freedom and constraint. The point is not whether there is a boundary. The point is what feedback the boundary produces.

A self-sustaining boundary moves some regulation inside the system. It may still use rules and sanctions, but these devices support a pattern participants increasingly carry themselves. An externally sustained

The Governability Gap

boundary holds behavior in place by adding pressure. It can be necessary for a time. It may be the only available response in crisis. But it should not be mistaken for restored viability.

The feedback test is simple. After a boundary is introduced, does the cost of external control fall, or does it rise?

If control costs fall, the boundary may be building the foundation. If they rise, the boundary is preserving form while drawing down the foundation.

That is a small distinction. It changes the diagnosis.

7 | Diagnosing Viability Thresholds

The accepted abstract for this paper describes the governability gap as a viability threshold. The question is how such a threshold can be diagnosed before collapse.

The proposed answer is not to look for one master indicator. Complex systems rarely give that kind of warning. The paper suggests a pattern instead.

First, the governance model becomes increasingly output-centered. It sees throughput, targets, performance, compliance, and reportable units more clearly than it sees maintenance and foundations.

Second, maintenance becomes chronically underprotected. Recovery, repair, care, prevention, mentoring, and informal coordination are treated as secondary because they do not produce clear output units.

Third, foundations are converted into proxies. Trust becomes a score. Meaning becomes engagement. Legitimacy becomes formal acceptance. Judgment becomes compliance with procedure.

Fourth, SCI rises. The system must spend more on control, incentives, verification, enforcement, symbolic cohesion, or procedural specification to maintain the same core output.

Fifth, boundary problems appear. The system either cannot build a boundary to protect foundation-preserving behavior, or it builds one that requires increasing external control.

This gives a compact diagnostic chain:

Visible output becomes dominant. Maintenance is undercounted. Foundations are proxied. Substitution costs rise. Boundary-constructability determines whether protection is possible. Boundary feedback shows whether the protection is rebuilding the foundation or replacing it. Viability depends on the result. Exhibit 1 summarizes the diagnostic chain from visible output to viability threshold.

Exhibit 1. Diagnostic Chain from Visible Output to Viability Threshold.

Stage	Diagnostic Question	Warning Signal
Output model	What does the system count as success?	Narrowing of attention to visible throughput
Maintenance	What keeps output possible?	Non-events, repair, recovery, and prevention disappear from the model
Foundations	What lowers coordination cost?	Trust, meaning, legitimacy, and judgment are replaced by proxies
SCI	What is the system paying to replace foundations?	Rising control, audit, incentives, enforcement, or procedural load
Boundary	Can foundation-preserving behavior be protected?	Defection is invisible or response is illegitimate
Boundary feedback	Does control become cheaper or more expensive?	Rising external control after boundary introduction
Viability	Can the system absorb variety without drawing down foundations?	Stable output with rising substitution costs

This is not a prediction machine. It is a way to ask better questions. Three propositions follow.

The Governability Gap

First, in output-centered governance regimes, substitution costs will tend to rise before visible output declines, especially where coordination depends on trust, legitimacy, meaning, or professional judgment. Second, high boundary-constructability will weaken the link between output pressure and SCI growth, because the system can protect foundation-preserving behaviour from short-term extraction.

Third, after a boundary is introduced, self-sustaining boundaries will be followed by declining external control costs, while externally sustained boundaries will be followed by rising control costs and weaker viability signals.

These propositions can be investigated through comparative case studies, organizational time-series data, process tracing, or mixed-method diagnostics. The measurement will be imperfect. That is expected. The point is to track whether the system increasingly pays for substitutes.

8 | Discussion: Design Under Complexity Constraints

The design implication is not measure less. It is more precise.

Governance must learn which invisibilities call for better visibility and which call for protection from direct measurement.

For Type One invisibility, better recognition is often needed. Maintenance should be built into organizational design. Workload models should include repair and recovery. Public systems should value prevention, not only response. Universities should recognize mentoring, program repair, and intellectual maintenance. Hospitals should protect coordination time and recovery capacity. Technical systems already know this lesson better: preventive maintenance is not wasted time. It is system capacity preserved.

For Type Two invisibility, the design response is different. Foundations should be studied, but not turned too quickly into direct targets. The safer route is indirect diagnosis, minimal critical specification, local judgment, and boundary design.

This connects to the design agenda named in the abstract: modularity, distributed sense-making, and adaptive variety alignment.

Modularity matters because local units can preserve context. They can hold more tacit knowledge than distant centers. They can see weak signals earlier. But modularity without boundary becomes fragmentation. The module must be inside a system of obligation.

Distributed sense-making matters because no centre can hold the full variety of the field. But distributed sense-making without transparency becomes blindness. Local units need access to the state of the system, not just freedom from command.

Adaptive variety alignment matters because governance should not simply add more control. It should align the form of control with the kind of variety faced. Some variety must be reduced. Some must be preserved. Some must be left to local adaptation. Some must be contained by boundary.

This is where many governance reforms go wrong. They add organization where the problem is depleted foundation. They add reporting where the problem is lost judgment. They add incentives where the problem is lost meaning. They add enforcement where the problem is legitimacy. Each intervention may make sense locally. Together they can raise SCI and move the system closer to a viability threshold.

The better question is: what is the cheapest form of coordination that can still hold the system's purpose? Sometimes the answer is a metric. Sometimes a rule. Sometimes a boundary. Sometimes a local forum.

Sometimes professional autonomy. Sometimes a hard constraint with freedom inside it.

There is no universal preference for looseness. Loose systems can fail. Small systems can become parochial and coercive. Trust can hide abuse. Local rules can protect insiders. Self-sustaining boundaries can sustain harmful patterns. The framework is functional, not moral. It asks what a governance form does to the system's capacity to absorb variety without drawing down its foundations.

A later normative analysis can ask whether the pattern should be preserved. That is a different question.

9 | Limitations

The first limitation is conceptual. Maintenance and foundations are analytically distinct, but in practice they overlap. Maintenance often protects foundations. Foundations make maintenance cheaper. The distinction is useful because the governance response differs, not because the world separates cleanly. The second limitation concerns SCI. It is a heuristic, not a validated index. It requires interpretation. A rising compliance cost may signal foundation depletion, but it may also reflect necessary correction of a previously unsafe or unjust system. An increase in enforcement may mean legitimacy has declined; it may also mean previous non-enforcement protected powerful actors. SCI does not decide this. It marks the need for diagnosis.

The third limitation is boundary definition. Boundary-constructability can be difficult to assess. Boundaries are formal and informal, visible and tacit. A system may have a legal boundary but no social boundary. It may have a strong internal identity and weak accountability. More empirical work is needed to specify the variables and weights.

The fourth limitation is scale. The paper uses organizational examples to clarify mechanisms. Extension to national, planetary, or digital platform governance needs additional work. Large-scale systems may require nested boundary structures rather than one boundary.

The fifth limitation concerns the accepted abstract's broad claim about contemporary governance failure. This paper narrows that claim. It does not argue that all governance failure follows from the mechanism described here. It argues that one recurrent form of failure can be diagnosed through the interaction of model, maintenance, foundation, substitution cost, and boundary.

The narrower claim is stronger.

10 | Conclusion

The governability gap is usually described as a mismatch between the complexity of the governed system and the capacity of the governing system. That is correct, but incomplete.

The governing system acts through a model. The model makes some things visible and others hard to see. Output becomes legible. Maintenance often becomes a non-event. Foundations become proxies. The regulator then receives clean signals from the layer it can count and weak signals from the layer that keeps the system viable.

The Substitution-Cost Indicator offers one way to detect the problem. It does not measure trust, meaning, legitimacy, or judgment directly. It asks what the system must increasingly pay to replace them. Rising control, audit, incentives, enforcement, symbolic cohesion, and procedural specification can be early signs that the system is holding visible output while drawing down its foundations.

Boundary-constructability then asks whether the foundation can be protected. Boundary feedback asks whether the protection works. If a boundary reduces the need for external control, it may be rebuilding the foundation. If it requires rising control, it is preserving form while increasing substitution cost.

This gives a practical reading of the viability threshold. A system approaches that threshold when it can maintain output only by raising the cost of coordination, while losing the foundations that once made coordination cheaper.

The final design lesson is modest. Governance of complex systems cannot improve simply by adding more organization. It must distinguish what should be measured, what should be maintained, what should be protected from direct measurement, and what kind of boundary can hold that protection without becoming another drain on the foundation.

That is where adaptive governance begins.

References

- Ashby, W. R. (1956). *An introduction to cybernetics*. Chapman & Hall.
- Beer, S. (1972). *Brain of the firm*. Allen Lane.
- Beer, S. (1979). *The heart of enterprise*. Wiley.
- Beer, S. (1985). *Diagnosing the system for organizations*. Wiley.

The Governability Gap

- Campbell, D. T. (1976). Assessing the impact of planned social change (Occasional Paper Series No. 8). The Public Affairs Center, Dartmouth College.
- Cherns, A. (1976). The principles of sociotechnical design. *Human Relations*, 29(8), 783–792.
- Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97.
<https://doi.org/10.1080/00207727008920220>
- Goodhart, C. A. E. (1975). Problems of monetary management: The U.K. experience. *Papers in Monetary Economics*, 1. Reserve Bank of Australia.
- Graham, S., & Thrift, N. (2007). Out of order: Understanding repair and maintenance. *Theory, Culture & Society*, 24(3), 1–25.
- Jackson, S. J. (2014). Rethinking repair. In T. Gillespie, P. J. Boczkowski, & K. A. Foot (Eds.), *Media technologies: Essays on communication, materiality, and society* (pp. 221–239). MIT Press.
- Kooiman, J. (2003). *Governing as governance*. SAGE.
- Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*. D. Reidel.
- Olson, M. (1965). *The logic of collective action: Public goods and the theory of groups*. Harvard University Press.
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.
- Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition have failed*. Yale University Press.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118.
- Simon, H. A. (1957). *Models of man: Social and rational*. Wiley.
- Star, S. L., & Strauss, A. (1999). Layers of silence, arenas of voice: The ecology of visible and invisible work. *Computer Supported Cooperative Work*, 8, 9–30.
- Strathern, M. (1997). Improving ratings: Audit in the British university system. *European Review*, 5(3), 305–321.
- von Foerster, H. (1981). *Observing systems*. Intersystems Publications.