The Dashboard Was Green

Why Ecological Governance Keeps Failing—and What Would Actually Work

A synthesis for the climate, ecological, and global governance community.

Björn Kenneth Holmström

June 2026

The Clouded Mirror · Reader's Guide

`https://bjornkennethholmstrom.org/syntheses/the-dashboard-was-green`

In 1992, the Northern cod fishery off the coast of Newfoundland collapsed. Forty thousand people lost their livelihoods. The ecosystem, once one of the most productive on earth, has not recovered. The collapse was not a failure of science or governance in the conventional sense. The stock assessments were rigorous. The quotas were based on the best available data. The scientists were competent, the regulators were diligent, and the fishers, for the most part, followed the rules. The dashboard showed sustainable stock levels until the moment it didn't. The collapse appeared sudden, unexpected, and inexplicable.

What happened is now understood as a failure of perception. The stock assessments tracked a single metric: total biomass, estimated from annual surveys. They did not track the spatial distribution of fish populations. They did not track the age structure of the stock. They did not track the health of the food web, the reproductive success of non-commercial species, or the slow ecological signals that indicated the ecosystem was approaching a regime shift. The governance system was optimising a one-dimensional model of a multi-dimensional reality. It authorised extraction at levels the model said were sustainable. The ecosystem could not sustain them. The dashboard was green. The ecosystem died. The two facts are causally connected.

This is not a story about fisheries management. It is a story about the structural condition that produces ecological collapse across every domain we govern. We call it the *Variety Gap*: the mismatch between the effective dimensionality of the ecosystems we must manage and the effective dimensionality of the observation channels through which we try to perceive them. The gap is growing. And the consequences of that growth are already visible in the converging crises of the twenty-first century.

The Commons as a Feedback Problem

The standard account of ecological collapse—the tragedy of the commons—locates the problem in incentives. Individuals over-extract because they capture the full benefit while sharing the cost of depletion. The solution, in this framing, is to change the incentives: privatise the resource so owners bear the full cost of their decisions, or impose state regulation to limit extraction below the collectively ruinous level.

This account is not wrong, but it is incomplete. The deeper problem is architectural before it is motivational. Individual extraction decisions made without feedback from the collective resource state constitute an open-loop system: an actuator with no sensor. Each user makes decisions based on their immediate circumstances, without coupling to the aggregate stock or the aggregate extraction of others. The collapse is caused not by greed but by the absence of a closed feedback loop between what each actor does and what the collective consequences are. Improving individual motivation in an open-loop system does not close the loop. Closing the loop requires observation.

And observation has structural prerequisites. Ashby's Law of Requisite Variety, a foundational theorem of cybernetics, states that a controller can only stabilise a system if its internal variety matches or exceeds the variety of the disturbances it faces. For ecological governance, the "controller" is the set of institutions, metrics, and decision processes that translate observations of the resource into extraction rules. Its variety is

the effective dimensionality of its observation channel: how many independent aspects of the ecosystem it can perceive and respond to. The ecosystem itself—a fishery, a forest, a watershed, a climate system—has enormous variety. It generates disturbances across multiple timescales simultaneously: fast stochastic shocks, medium seasonal cycles, slow decadal shifts in carrying capacity. Governing it requires an observation channel that covers all three frequency bands.

The Variety Gap is the mismatch between the dimensionality of the ecosystem and the dimensionality of the governance system's perception. When the gap is large, the unobserved variety appears as uncontrolled variance in the outcomes. The system authorises extraction at levels that the full complexity of the ecosystem cannot sustain. The collapse occurs along the unmeasured dimensions. It appears sudden and inexplicable because the dashboard was only watching a fraction of what mattered.

Why State Management Often Makes Things Worse

The most counterintuitive finding to emerge from this analysis is that state management—the standard post-Hardin response of centralised regulation, annual surveys, and scientifically-determined quotas—can perform worse than open access. The mechanism is precise. State management combines two properties: high observation latency (annual surveys) and single-dimension aggregation (total biomass only). In a declining resource, the quota issued at time t is calibrated to the stock observed at time t minus twelve months. If the stock is declining, last year's observation is higher than this year's reality. The quota authorises extraction at a level that was sustainable for last year's conditions but is above sustainable yield for this year's. Each annual cycle extracts slightly more than the current stock can regenerate. The intervention designed to prevent overharvesting becomes the mechanism of it.

Open access, by contrast, has no lag. Each user responds to local conditions immediately. The feedback is weak and uncoordinated, but it is immediate. In a multi-scale disturbance environment, immediate weak feedback can outperform delayed strong authorisation. This is not an argument against governance. It is an argument about the architectural prerequisites for governance to be beneficial. Authority without observation quality is a blindfolded controller applying inputs to a system it cannot see—sometimes by chance stabilising it, sometimes accelerating its collapse. The prerequisite for effective ecological governance is not enforcement capacity or institutional quality. It is perceptual capacity. And that capacity is determined by where the observer stands.

The Ostrom Jump and the Slow Variable Jump

Elinor Ostrom's Nobel Prize-winning work demonstrated that communities around the world successfully manage shared resources without either privatisation or state regulation. Her eight design principles—clearly defined boundaries, rules matched to local conditions, collective-choice arrangements, monitoring, graduated

sanctions, conflict resolution mechanisms, recognition by external authorities, and nested governance for larger systems—describe the institutional features of enduring commons governance. What the Variety Gap framework adds is an explanation for *why* these principles work.

A governance simulation developed as part of the research programme compared five architectures for renewable resource management under identical disturbance conditions: open access, state management, market mechanisms, Ostrom-style community commons, and indigenous bioregional governance. The community commons architecture, with monthly monitoring, multi-dimensional observation (stock level, spatial distribution, social pressure signals), and graduated sanctions, reduced collapse risk from over eighty percent to roughly thirty percent. The performance jump was not attributable to stronger values or more committed actors. It was attributable to the expansion of observation dimensionality from one (the market price signal, or the annual aggregate stock) to three. Ostrom's design principles work because they open additional observation channels. The community as a distributed sensor network perceives dimensions—spatial variation, social dynamics, early warning signals from specific locations—that are structurally invisible to a central authority observing an annual aggregate.

But the community commons architecture still has a critical blind spot. It operates with monthly monitoring over years to decades. It cannot perceive the slow decadal trends in carrying capacity that are driven by climate change, land use transformation, or long-run ecological dynamics. The slow variable is invisible to any governance system whose observation window is shorter than the trend itself. Only the indigenous bioregional architecture—which adds access to the slow ecological signal through multi-generational observation records, seasonal phenological indicators, and species co-occurrence signals—reduced collapse risk to under four percent. It was the only architecture whose trajectory showed any correspondence to the true carrying capacity trend. All other architectures discovered the slow decline retrospectively, as crisis.

This is the *slow variable jump*: the performance improvement that comes from adding access to the decadal and centennial dimensions of ecosystem dynamics. It is the dimension that traditional ecological knowledge provides, and that no monitoring programme operating on administrative timescales can replicate. A community that has managed a fishery for five hundred years has five hundred years of slow-variable signal embedded in its governance protocols. The seasonal restrictions encode observations about what happens when those restrictions are violated. The taboos around certain species reflect accumulated knowledge about ecological thresholds. The knowledge is not mystical. It is the output of a continuous, multi-dimensional observation process that modern science has simply not existed long enough to conduct.

Indigenous Sovereignty as an Engineering Requirement

The recognition of indigenous land rights and resource sovereignty has been framed primarily in terms of historical justice—the rectification of colonial dispossession—and cultural rights—the protection of ways of life that depend on land access. These framings are legitimate and important. They are also incomplete.

The Variety Gap framework adds a third framing that does not depend on historical or cultural arguments: indigenous resource sovereignty is an *engineering requirement* for effective ecological governance. Communities with multi-generational embeddedness in an ecosystem have accumulated, through continuous observation across generations, the observation dimensionality required to govern that ecosystem across all relevant disturbance timescales. Displacing this governance with external administration that has the properties of state management—high latency, single-dimension aggregation, no access to the slow variable signal—does not simply change who holds authority. It destroys the observation system that was performing the slow-variable governance function. The institutional knowledge embedded in the community's protocols, seasonal practices, and land-use traditions is not transferred when authority changes hands. It is lost. And it cannot be replaced by any amount of well-funded monitoring infrastructure operating on administrative timescales.

A monitoring programme that has existed for ten years does not have access to the slow variable signals that a community observation record accumulated over centuries provides. The effective governance capacity that was lost through dispossession cannot be reconstructed by improving the technology available to external managers. It can only be accessed through governance systems that have been present and observing continuously across the relevant timescales—which, for decadal and longer ecological dynamics, means governance systems embedded in the landscape across generations.

This is not a claim that all indigenous governance systems always achieve superior outcomes, or that traditional practice is beyond critique or adaptation. It is the structural observation that governance systems with the requisite variety properties that long-run ecological embeddedness provides will outperform governance systems without those properties. The primary existing source of such properties for most of the world's complex renewable resources is the communities that have governed them across generations. Recognising their sovereignty is not an act of cultural generosity. It is a prerequisite for managing those resources at all.

The Planetary Variety Gap

Climate change is the commons problem at planetary scale: a shared atmospheric resource degraded by the aggregate of distributed extraction decisions, governed by institutions that lack the observation dimensionality and temporal resolution to manage it effectively. The slow variable problem that the cod fishery simulation identifies has a direct analogue in climate governance. The trend in atmospheric carbon concentration has been accumulating for two centuries. The governance response—meaningful international agreements—arrived only after the trend had become large enough to exceed the noise floor of short-window political attention. Like every architecture in the simulation except the one with access to the slow ecological signal, climate governance discovered the slow variable in retrospect, as a crisis rather than as an early signal.

The communities most likely to possess the slow ecological signal for climate impacts—communities observing glaciers, seasonal timing, species range shifts, and weather pattern changes across generations—are often the same communities whose resource sovereignty is most contested. Indigenous Arctic communities observed the slow warming signal in permafrost, sea ice, and wildlife behaviour before instrumental records confirmed it. Pastoral communities observed shifting rainfall patterns before modelling frameworks could predict them. Fishing communities observed regime shifts in marine ecosystems before stock assessments caught up. The epistemic value of these observations is not sentimental. It is the slow variable signal dimension that the planetary governance system most urgently needs and structurally lacks.

The Variety Gap at the planetary scale is the mismatch between the dimensionality of the Earth system—climate, biosphere, cryosphere, hydrosphere, and all their nonlinear interactions—and the dimensionality of the international institutions tasked with governing it. The gap is not merely large. It is widening. The acceleration asymmetry—the gap between the rate at which the Earth system generates novel disturbance dimensions and the rate at which governance institutions expand to perceive them—is the defining structural condition of the Anthropocene.

Closing that gap, even partially, requires an architecture that embeds observation within the systems being governed, matches response speed to disturbance timescale at each level, protects feedback loops from capture by the actors they evaluate, and maintains a permanent capacity to add new dimensions of perception as the environment continues to change. The building blocks of that architecture exist. Bioregional governance units defined by ecological boundaries rather than administrative ones. Commons trusts that hold renewable resources in stewardship rather than ownership. Indigenous wisdom councils that provide access to the slow ecological signal. Nested, polycentric coordination that matches each governance function to the scale at which it can actually be performed.

These are not utopian proposals. They are structural responses to a structural diagnosis. The fragments are already here. The question is whether they will be connected—or whether the planetary governance system will continue to optimise a narrow set of metrics while the excluded dimensions of the Earth system force a reckoning that no existing institution can survive.

The Invitation

The ecological governance community—climate scientists, conservation biologists, environmental economists, indigenous rights advocates, bioregional organisers—is already working on the dimensions that the standard governance architecture excludes. It tracks what the dashboard leaves invisible. It advocates for the communities whose observation capacity the state management model destroys. It builds the alternative institutions—land trusts, watershed councils, community-managed marine areas—that demonstrate higher-dimensional governance in practice.

What the community lacks is a formal language for why these alternatives work, a measurement framework for distinguishing genuine ecological governance from symbolic sustainability, and a strategic logic for connecting the fragments into an architecture that can scale. The Variety Gap framework offers all three. It provides the formal rationale: observation dimensionality determines governance performance, and the slow variable dimension is structurally inaccessible to external managers operating on administrative timescales. It provides the measurement instrument: a set of parameters—observation dimensionality, response latency, signal fidelity, immune permeability—that can estimate whether a governance reform is actually expanding perceptual capacity or merely rearranging it. And it provides the strategic logic: the central architecture cannot perceive the need for its own transformation. The only viable path is to build the alternative at the periphery, protect it from capture, and let the evidence accumulate until the cost of maintaining the current architecture becomes undeniable.

The invitation is to test this framework in practice. The first step is a *Variety Gap audit* of an ecological governance system: a fishery, a forest, a watershed, a conservation area. The audit would estimate the current observation dimensionality, identify the excluded dimensions most causally relevant to long-run ecological viability, and assess whether the governance architecture is tracking the slow variables that determine whether the system is approaching a regime shift. The audit does not require new technology. It can be conducted with existing ecological data, structured expert elicitation, and the open-source measurement protocol that accompanies the framework.

The cod fishery collapsed because the dashboard was green while the ecosystem died. The climate is changing because the dashboard—the emissions pledges, the GDP growth rates, the technology deployment curves—remains green while the Earth system approaches thresholds that the dashboard was never designed to see. The question is not whether we need better metrics. We do. The question is whether we can build the institutional capacity to keep adding metrics, indefinitely, as the planet continues to generate dimensions that no existing dashboard can capture.

The fragments of that capacity are already here. They have been here for centuries, in the communities that never lost them. The work of connecting them—of protecting them, of learning from them, of scaling their demonstrated performance into an architecture that can meet the acceleration asymmetry—is the defining governance challenge of our time.

The measurement begins here. It does not end here. The invitation is open.