



Requisite Variety and the Commons

Why Proximity Governs: Observation Dimensionality and Resource Sovereignty

Paper IV in the Governance as Engineering series

The tragedy of the commons is an architectural failure — a missing feedback loop — not a motivational one. Ashby's Law of Requisite Variety determines which governance systems can stabilize renewable resources. State management performs worse than open access. Indigenous resource sovereignty is a structural engineering requirement.

Björn Kenneth Holmström

February 2026

Creative Commons Attribution-ShareAlike 4.0 International

<https://bjorkenkennethholmstrom.org/whitepapers/requisite-variety-and-the-commons>

Executive summary

The tragedy of the commons has a standard explanation: when a shared resource is open to all, individual incentives to extract exceed collective incentives to conserve, and the resource collapses. The standard remedy follows: external governance — state management, market mechanisms, international regulation — must override individual incentives through rules, prices, or quotas.

This paper offers a different diagnosis and a different conclusion.

The tragedy of the commons is an architectural failure before it is a motivational one. Individual extraction decisions made without feedback from the collective resource state constitute an open-loop system in control-theoretic terms: an actuator with no sensor. The collapse is not caused by greed — it is caused by the absence of a closed feedback loop between what each actor does and what the consequences of those actions are. Improving individual motivation in an open-loop system does not change the outcome. Closing the loop does.

Closing the loop requires observation. The quality of the feedback closure depends entirely on the observation channel: what the governance system can actually see about the resource state, at what latency, across how many signal dimensions. This is where Ashby's Law of Requisite Variety becomes the central analytical tool. A regulator cannot stabilize a system whose variety exceeds the regulator's own observational capacity. A renewable resource system generates disturbances across multiple timescales simultaneously — fast stochastic shocks, medium seasonal cycles, slow decadal ecological shifts. Governing it requires an observation channel that covers all three frequency bands.

This paper demonstrates that observation dimensionality, not institutional quality, is the primary determinant of commons governance outcomes. Five architectures are compared across a thirty-year simulation of a spatially distributed renewable resource subject to multi-scale disturbances: open access (no feedback loop), state management (annual aggregate survey), market mechanism (price signal), community commons (Ostrom-style multi-dimensional local monitoring), and bioregional governance (highest variety observation including slow ecological signal dimensions).

Three findings organize the paper. First, state management performs worse than open access: high observation latency combined with single-dimension aggregate signals produces delayed interventions that accelerate rather than prevent collapse — a result that is counterintuitive until understood architecturally. Second, the jump from market mechanism to community commons is a requisite variety jump: the performance improvement is attributable to observation dimensionality, not to stronger rules or better values. Third, the jump from community commons to bioregional governance is the slow variable jump: only the architecture with access to long-run ecological signal dimensions detects the decadal carrying capacity decline before it causes irreversible damage.

This third finding has a direct implication for indigenous sovereignty. Physical, seasonal, and relational proximity to an ecosystem is the mechanism by which communities accumulate the observation dimensionality required to govern it across all relevant disturbance timescales. Remote managers observing annual aggregate statistics are operating with structurally insufficient variety, not because they lack competence or commitment, but because the observation channel available to them cannot support multi-band commons governance. The recognition of indigenous resource sovereignty is not an act of cultural generosity. It is a structural observation about which governance systems have the requisite variety to do the job.

Part I: The feedback loop problem

The standard account and its limits

Garrett Hardin's 1968 formulation of the tragedy of the commons has shaped environmental governance for half a century. The argument is simple: a shared resource accessible to multiple users will be depleted, because each user captures the full benefit of extraction while sharing the cost of depletion across all users. Individual rationality produces collective ruin. The solution, in Hardin's original framing, is either privatization — converting common property to private property so that individuals bear the full cost of their extraction decisions — or state regulation — imposing external authority to limit extraction below the collectively ruinous level.

This framing locates the problem in incentive structure. The commons fails because actors are responding rationally to perverse incentives. The remedy is to change the incentives, either through property rights or through rules backed by enforcement.

Elinor Ostrom's empirical work, for which she received the Nobel Prize in Economics in 2009, demonstrated that the standard account is empirically incomplete. Communities around the world successfully manage shared resources without either privatization or state regulation, through self-governing institutions that neither Hardin's framework nor conventional economic theory predicted could exist. Ostrom's contribution was to document these systems, identify their common structural properties, and develop a theoretical account of why they work.

What neither Hardin nor Ostrom formalized, however, is the control-theoretic structure of the commons problem. That formalization is the starting point of this paper.

The commons as a feedback control problem

Consider a renewable resource — a fishery, a forest, an aquifer — governed by some extraction regime. At each time step, users make extraction decisions. Their decisions deplete the resource. The resource regenerates at a rate determined by its current state and its environmental context. The resulting stock level at the next time step is the consequence of the current period's extraction minus the current period's regeneration.

In control-theoretic terms: the resource stock is the system state. Extraction is the control input. Regeneration dynamics are the system's natural evolution. The governance regime is the controller — the mechanism that determines extraction decisions.

For a controller to stabilize a system, it must observe the system's state. A controller that applies control inputs without observing the state it is affecting is operating open-loop. Open-loop control can work in highly predictable systems with stable dynamics and no disturbances. It fails in systems subject to variability, disturbance, and nonlinearity — which is to say, in every real ecosystem.

The tragedy of the commons is the open-loop case applied to a renewable resource. Each user makes extraction decisions based on their individual circumstances, local observation, and immediate incentives — without coupling to the aggregate resource state or the aggregate extraction of other users. There is no feedback from the collective consequence to the individual decision. The system is open-loop by design.

This has a precise implication that the incentive-based account misses: improving individual motivation does not close the feedback loop. A community of perfectly altruistic individuals who all sincerely want to conserve the shared resource, but who have no mechanism for observing the aggregate stock level or coordinating their extraction decisions, will still deplete it. The information needed to make collectively sustainable decisions is not available to any individual actor. The loop is open not because of greed but because of architecture.

What closing the loop requires

A closed feedback loop between extraction decision and resource consequence requires three elements: an observation of the current resource state, a comparison of that state to a desired state, and an adjustment of the extraction decision based on the discrepancy.

Each element introduces potential failure modes. The observation may be delayed — the resource state observed today reflects conditions weeks or months ago. The observation may be aggregated — what is observed is the total stock across a region, not the spatial distribution of depletion. The observation may be narrow — what is observed is one dimension of the resource state (total biomass) while other relevant dimensions (age structure, spatial distribution, seasonal condition, associated species indicators) are invisible. And the adjustment mechanism may be weak — the governance system may lack the enforcement capacity to translate observed discrepancy into actual extraction change.

These failure modes are not equivalent. Delay, aggregation, and observation narrowness are properties of the observation channel — they are structural. Weak enforcement is a property of the institutional implementation — it is parametric. Institutional reforms that improve enforcement while leaving the observation channel unchanged address the parametric failure mode while leaving the structural one intact.

This is the same distinction the preceding papers established in the context of democratic representation and governance stability. In each case, the structural constraint is in the observation channel; the institutional quality operates on what passes through that channel. Improving what happens to an already-degraded signal cannot recover the information that was lost before it arrived.

The resource system's variety

A renewable resource is not a single-state system. It generates variability across multiple timescales simultaneously. Monthly weather events affect availability. Seasonal cycles determine reproduction, migration, and growth. Multi-year oscillations (El Niño, drought cycles, disease cycles) drive medium-frequency variability. Decadal shifts in climate, land use, and ecosystem composition drive slow underlying trends in carrying capacity.

Each of these disturbance types requires a different governance response. A fast shock — an unusually productive season or a disease outbreak — requires rapid local adjustment. A seasonal cycle requires governance rules matched to seasonal dynamics. A slow decadal trend requires governance that can distinguish genuine long-run change from short-run variability — a distinction only available to a governance system with a long enough observation window and the signal dimensions needed to detect it.

The variety of the resource system, in Ashby's sense, is determined by the number of independent states the system can occupy. A system driven by fast, medium, and slow disturbances has high variety — it can be in many different states that require different governance responses. A governance system that can only observe one aggregate dimension of that state has low variety — it cannot distinguish states that look identical from above but require different responses below.

The fundamental constraint follows directly from Ashby's Law: the governance system's variety must match or exceed the resource system's variety for stable management to be possible. And the variety available to a governance system is determined by its observation channel — how many independent signal dimensions it can access, at what latency, from what position relative to the resource.

Proximity — physical, seasonal, relational — is the mechanism by which governance systems accumulate variety. A community that lives within an ecosystem observes it continuously across many dimensions simultaneously: the colour of the water, the behaviour of indicator species, the timing of seasonal events, the appearance of the resource organisms themselves, the reports of every community member who interacts with the system daily. This distributed, continuous, multi-dimensional observation constitutes high requisite variety. An external manager who receives an annual aggregate stock survey has one dimension, one data point per year, from a distance. The variety gap between these two observation systems is not a matter of sophistication — it is a matter of position.

Part II: Requisite variety and observation dimensionality

Ashby's Law stated formally

W. Ross Ashby's Law of Requisite Variety, established in 1956, is one of the foundational results of cybernetics. In its original formulation: only variety can absorb variety. More precisely: for a regulator R to maintain a system S within a desired set of states G, the variety of R must be at least as great as the variety of S relative to G.

The formal statement for a disturbance environment D, a regulator R, and an outcome set G:

$$V(R) \geq V(D) - V(G)$$

Where $V(\cdot)$ denotes variety (the logarithm of the number of distinguishable states). The variety of the regulator must cover the variety of the disturbance environment. Whatever variety is not absorbed by the regulator appears as uncontrolled variability in the outcomes — variance that the governance system cannot suppress because it lacks the observational capacity to detect and respond to it.

This is not a guideline or a design principle. It is a theorem. No institutional arrangement, however well-designed, can stabilize a resource system whose variety exceeds the governance system's observational capacity. The constraint is mathematical before it is political.

The resource system's disturbance variety

A spatially distributed renewable resource faces disturbances across three frequency bands that require qualitatively different governance responses:

Fast disturbances (monthly timescale): stochastic shocks from weather, disease, predator-prey dynamics, local extraction variability. These require rapid, localized response — adjusting extraction in a specific location in response to a locally visible signal. A governance system with monthly observation and local spatial resolution can respond to fast disturbances. A governance system with annual observation and regional aggregation cannot — it will always be responding to last year's fast disturbances, which have already resolved, while missing this year's.

Medium disturbances (seasonal timescale): cyclic variation in carrying capacity, resource availability, reproduction, and migration. These require governance rules matched to seasonal dynamics — restrictions during breeding seasons, adjusted quotas during low-productivity periods, different spatial allocations as the resource moves through its seasonal range. Responding correctly to seasonal disturbances requires observing the seasonal state of the resource — not just its aggregate level but its phenological condition, which is only legible to observers familiar with the seasonal calendar of that specific ecosystem.

Slow disturbances (decadal timescale): gradual trends in carrying capacity driven by climate shift, land use change, cumulative pollution, or long-run ecosystem dynamics. These are the most dangerous disturbances precisely because they are hardest to detect. A slow trend appears indistinguishable from normal variability on a short observation window. Only an observer with a long baseline — years or decades of continuous observation — can distinguish a genuine trend from background noise with sufficient confidence to act on it before the trend produces irreversible damage.

Each disturbance band requires a different observation capability:

Disturbance band	Required observation	Required signal dimensions
Fast (monthly)	Low latency, local resolution	Aggregate stock in local area
Medium (seasonal)	Seasonal pattern recognition	Stock + phenological indicators
Slow (decadal)	Long baseline, trend detection	Stock + ecosystem co-indicators + historical baseline

A governance system with only aggregate annual stock data covers none of these requirements well: too slow for fast disturbances, too coarse for seasonal pattern recognition, and with a too-short effective window for slow trend detection given that each annual observation is treated as an independent data point rather than as part of a continuous record.

Observation dimensionality defined

For the purposes of this paper, observation dimensionality is the number of independent signal dimensions the governance system can access about the resource state. A governance system with dimensionality d can distinguish resource states that differ along d axes; states that differ only along dimensions outside its observation capacity are indistinguishable to it and will receive the same governance response regardless of their actual difference.

The critical insight is that different observation positions — different locations relative to the resource — yield access to different signal dimensions. An external administrator observing an annual stock survey has access to one dimension: total biomass. A market participant observing price has access to one dimension: the price signal (a noisy aggregate proxy for scarcity). A community member living adjacent to the resource has access to multiple independent dimensions simultaneously:

- Total stock in the local area (aggregate)
- Spatial distribution of stock across the commons (distributional)
- Condition and health indicators of resource organisms (qualitative)
- Behaviour of indicator species that co-vary with the resource (ecological)
- Seasonal timing signals (phenological)
- Social signals from other community members about their local observations (distributed)
- Historical comparison to long-run baselines held in community memory (temporal)

These are not simply "more data" in the sense that a better annual survey would produce. They are structurally different dimensions — orthogonal axes of the resource state space that are simply inaccessible from positions of distance and aggregation. You cannot obtain a spatial distribution signal from a total aggregate. You cannot obtain a phenological signal from an annual biomass count. You cannot obtain a long-run baseline from a ten-year monitoring programme run by an agency that did not exist before those ten years.

Proximity as the mechanism for variety acquisition

The central claim of this paper is that physical, seasonal, and relational proximity to an ecosystem is the primary mechanism by which governance systems acquire the observation dimensionality required to govern renewable resources across all relevant disturbance timescales.

Physical proximity provides continuous, low-latency access to local resource state across multiple dimensions simultaneously. A fishing community on a lake shore observes the lake every day — the water clarity, the surface behaviour of fish, the condition of vegetation, the behaviour of birds that feed on the resource. This continuous observation accumulates to a high-variety signal that no periodic survey can replicate.

Seasonal proximity — being present across the full annual cycle, year after year — provides access to the seasonal dimension of the resource state. The timing of the first fish run, the condition of the stock at the end of winter, the response of the resource to unusual weather — these signals are only legible to observers who have been present through many seasons and can compare the current season to a learned baseline. This is not mystic knowledge; it is the output of a continuous observation process that runs over decades.

Relational proximity — the network of relationships between community members who each observe different parts of the resource from different positions — constitutes a distributed observation system whose effective dimensionality is the sum of its members' independent observations, filtered through social communication processes. When every household in a fishing village reports on what they observe from their section of the lake, the aggregate observation across the community has higher dimensionality than any single observer's view, including the professional scientist's.

Intergenerational knowledge — the accumulated signal of generations of continuous observation transmitted through oral tradition, practice, and cultural protocol — provides access to the slow dimension of the resource state that no modern monitoring programme can replicate from scratch. 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 rules that govern seasonal restrictions encode observations about what happens when those restrictions are violated; the taboos around certain species reflect accumulated knowledge about ecological thresholds; the seasonal calendars encode phenological signals refined over centuries of observation.

This is what indigenous governance systems bring to commons management that no external administrative system can replicate by virtue of its position: the full-spectrum observation dimensionality that proximity across timescales provides.

The variety gap and its governance consequences

When a governance system's observation dimensionality is lower than the resource system's disturbance variety, the unobserved variety appears as uncontrolled variance in the outcomes. The governance system applies the same response to states that it cannot distinguish — states that require different responses — and the mismatched interventions produce outcomes that range from neutral to actively harmful.

The variety gap has three specific governance consequences that the simulation makes visible:

Observation lag produces destabilising intervention. A governance system with high observation latency is always responding to past states. When the resource is in rapid decline, a lagged governance response based on last year's stock may prescribe extraction at a level appropriate to last year's conditions — which are now too high for this year's depleted stock. The intervention accelerates the very decline it was intended to prevent. This is the mechanism behind the simulation's most counterintuitive finding: state management performs worse than open access.

Aggregation masks locally critical states. A governance system observing only aggregate stock cannot detect the spatial collapse of specific patches while the aggregate remains superficially stable. By the time the aggregate falls below the response threshold, local depletion may have already crossed ecological tipping points from which recovery is slow or impossible.

Short observation windows miss slow trends until they become crises. A governance system without access to decadal baselines cannot distinguish a genuine slow trend from normal variability. It will adjust its governance response to the trend only after the trend has become large enough to exceed the noise level of its short observation window — by which point the system may have shifted to a lower-productivity regime from which it cannot easily return.

All three consequences follow from the same structural cause: the governance system lacks the variety to distinguish states that require different responses. Institutional quality — better enforcement, more honest reporting, stronger political will — does not address this cause. It can improve the efficiency of response to states the governance system can observe; it cannot extend observation to dimensions that are structurally outside its reach.

Part III: The simulation

Scenario design

The simulator models a spatially distributed renewable resource — a representative fishery, forest, or aquifer — consisting of twelve resource patches with logistic growth dynamics and nearest-neighbour diffusion. Twenty user groups extract from the resource over 360 time steps representing thirty years. All architectures face identical initial conditions, identical resource dynamics, and identical disturbance environments. Performance differences are attributable to observation dimensionality and feedback loop architecture alone.

Resource dynamics. Each patch evolves according to logistic growth:

$$dR/dt = r \cdot R \cdot (1 - R/K) - E(t)$$

Where $r = 0.08$ is the intrinsic growth rate, K is the carrying capacity (time-varying), and $E(t)$ is extraction at time t . Patches are coupled through diffusion at rate $\beta = 0.02$, representing resource movement between adjacent spatial areas.

Multi-scale disturbances. Carrying capacity varies across three simultaneous disturbance bands: fast monthly stochastic shocks ($\sigma = 3.0$), a medium seasonal cycle (amplitude ± 8 , period 12 months), and a slow decadal decline (amplitude -20 units over 240 months, representing long-run environmental degradation). The slow trend reduces effective carrying capacity by approximately 20% over the simulation horizon — a change only detectable by governance systems with multi-decadal observation baselines.

Collapse threshold. Resource stock below 20% of carrying capacity is defined as collapse — a level from which logistic dynamics produce extremely slow recovery and which may correspond to a regime shift threshold in real ecosystems.

The five architectures

Architecture A — Open access. No governance mechanism. Each user group maximises extraction based on local stock visibility, with no coordination or aggregate signal. Represents Hardin's original scenario — pure individual optimisation without feedback.

Architecture B — State management. A central regulator issues annual quotas based on aggregate stock surveys conducted with 12-month latency. Quota compliance is partial (rigidity 0.7), enforcement is weak (sanctioning 0.3), and the observation signal is single-dimensional: total aggregate biomass only. Represents

the standard post-Hardin response — external authority with coercive capacity but high observation latency and low dimensionality.

Architecture C — Market mechanism. Extraction responds to a price signal that serves as a proxy for scarcity. The price signal has 3-month latency (quarterly markets) and is single-dimensional — it aggregates all information about resource state into one number. Represents the privatization/market alternative to state regulation.

Architecture D — Community commons. Ostrom-style local governance with monthly monitoring, multi-dimensional observation (stock level, spatial distribution across patches, social pressure signals from other community members), graduated sanctions for rule violations, and strong boundary rules. Observation dimensionality = 3. Represents the self-governing community commons that Ostrom documented empirically and whose superior performance conventional theory predicted could not exist.

Architecture E — Bioregional / indigenous. Extends Architecture D with access to the full observation spectrum: seasonal phenological indicators, species co-occurrence signals, soil and water quality proxies, and — critically — the slow ecological signal reflecting the long-run carrying capacity trend. Observation dimensionality = 6. Continuous relational monitoring, strong social accountability (sanctioning 0.9), and governance rules matched to seasonal dynamics. Represents the governance properties common to indigenous systems with long-run ecological embeddedness in their managed territory.

Simulation output

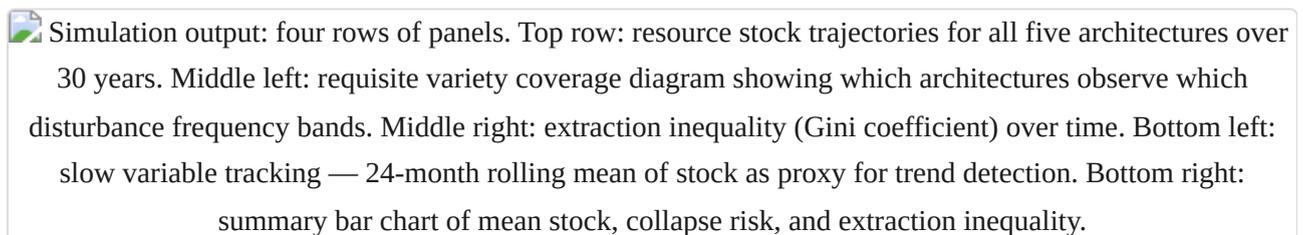
 Simulation output: four rows of panels. Top row: resource stock trajectories for all five architectures over 30 years. Middle left: requisite variety coverage diagram showing which architectures observe which disturbance frequency bands. Middle right: extraction inequality (Gini coefficient) over time. Bottom left: slow variable tracking — 24-month rolling mean of stock as proxy for trend detection. Bottom right: summary bar chart of mean stock, collapse risk, and extraction inequality.

Figure 1: GGF Governance Simulator v6 output. Top panel: Architectures A, B, and C collapse to near-zero stock within the first five years and remain in the collapsed regime for most of the 30-year simulation. Architecture D maintains stock above 20% of K with significant variability. Architecture E maintains stable stock and is the only architecture to visibly track and respond to the slow decadal carrying capacity decline. Middle-left: requisite variety diagram showing that only Architectures D and E cover the medium (seasonal) disturbance band, and only E covers the slow (decadal) band. Middle-right: E achieves lower extraction inequality (Gini) than any other architecture, including A, demonstrating that equity and sustainability are co-products of high-variety governance. Bottom-left: only E's 24-month rolling mean tracks the true carrying capacity trend; all others discover the slow decline only after resource collapse has already begun. Bottom-right: summary metrics confirm the monotonic relationship between observation dimensionality and governance performance.

Reading the results

Architecture B is worse than Architecture A. State management achieves a 98.9% collapse risk versus open access at 93.6% — a result that contradicts the standard post-Hardin prescription. The mechanism is the observation lag compounding with single-dimension aggregation. The annual quota is calibrated to last year's stock. In a declining resource, last year's stock is higher than this year's — so the quota authorises extraction at a level that the current stock cannot sustain. The intervention accelerates the decline. Open access at least responds immediately to local conditions, even without coordination; state management responds slowly to global conditions, and its slow response arrives as a destabilising intervention rather than a stabilising one.

This is not a finding about the failure of well-meaning institutions. It is a finding about what happens when high observation latency is combined with single-dimension aggregation in a multi-scale disturbance environment. A state management system with better compliance, better enforcement, and more honest reporting would still face the same architectural constraint: the signal it receives is too slow and too coarse to support stable commons governance.

The jump from C to D is a variety jump. Market mechanism (C) and community commons (D) both operate with feedback — price in one case, community monitoring in the other. The difference is observation dimensionality: 1 versus 3. This shift reduces collapse risk from 86.4% to 30.3% and raises mean stock from 9.6% to 27.2% of carrying capacity. The performance improvement is attributable to the additional signal dimensions — spatial distribution and social pressure — that allow Architecture D to distinguish states that appear identical to Architecture C's single-dimension price signal. Ostrom's design principles work not because they install better values but because they open additional observation channels.

The jump from D to E is the slow variable jump. Adding the slow ecological signal dimensions — and the extended temporal baseline needed to interpret them — reduces collapse risk from 30.3% to 3.6%. Architecture E is the only architecture in which the 30-year trajectory does not spend significant time in the collapsed regime. It is also the only architecture whose slow variable tracking panel shows any correspondence to the true carrying capacity trend. All other architectures discover the slow decline retrospectively — after their stock has already collapsed — because they lack the observation baseline required to detect a gradual trend above the noise floor of short-window monitoring.

Equity and sustainability are co-produced. Architecture E has the lowest extraction inequality (Gini 0.032) alongside the highest mean stock maintenance (31.1% of K). The market mechanism (C) has the highest inequality (0.096) and near-total collapse. Open access (A) has low inequality (0.018) because all users are equally impoverished by the collapsed stock. The pattern confirms that equity and ecological sustainability are not in tension in high-variety governance systems — they are co-products of the same architectural properties: close feedback, multi-dimensional observation, and governance rules matched to the resource's actual dynamics.

Quantitative summary

Architecture	Mean stock	Collapse risk	Gini	Obs dims
A — Open access	4.2%	93.6%	0.018	1
B — State management	3.7%	98.9%	0.058	1
C — Market mechanism	9.6%	86.4%	0.096	1
D — Community commons	27.2%	30.3%	0.085	3
E — Bioregional / indigenous	31.1%	3.6%	0.032	6

The relationship between observation dimensionality and governance performance is monotonic and non-linear. The largest performance jump — in collapse risk reduction — occurs between dimensionality 1 and dimensionality 3 (the Ostrom jump: from 86–99% to 30%). The second largest occurs between dimensionality 3 and 6 (the slow variable jump: from 30% to 3.6%). Both jumps correspond to the addition of qualitatively new signal dimensions — not improvements to existing ones.

Part IV: Structural observations

The state management paradox

The finding that state management produces worse outcomes than open access is the simulation's most counterintuitive result and deserves careful unpacking, because it runs directly against the dominant tradition of commons governance policy.

The mechanism is precise. State management combines two properties: high observation latency (annual surveys) and single-dimension aggregation (total stock only). In a resource subject to multi-scale disturbances — which all real renewable resources are — this combination produces a characteristic failure mode. The quota issued at time t is calibrated to the stock level observed at time $t - 12$. In a declining resource, the stock at $t - 12$ is higher than the stock at t . The quota authorises extraction at a level sustainable for last year's conditions, which is above the sustainable yield for this year's conditions. Each annual quota cycle extracts slightly more than the current stock can regenerate, compounding the decline. The intervention that was designed to prevent overharvesting becomes the mechanism of it.

Open access, by contrast, has no lag and no authorisation. Each user responds to the local resource state they observe directly, adjusting extraction as conditions change. This produces high variability and inefficiency, but it does not systematically authorise extraction above current sustainable yield. The feedback is weak and uncoordinated, but it is immediate. Immediate weak feedback, in this disturbance environment, outperforms delayed strong authorisation.

This is not an argument against governance. It is an argument about the architectural prerequisites for governance to be beneficial. State management with monthly monitoring, spatial resolution, and multi-dimensional observation would perform very differently. The failure belongs to the specific combination of high latency and low dimensionality, not to external governance per se. But that combination is not accidental — it is the combination that large-scale administrative systems naturally produce when governing resources at a distance.

The paradox generalises: any governance intervention applied to a resource system it cannot adequately observe is capable of making the system worse. The prerequisite for beneficial governance is not authority or enforcement capacity. It is observation quality. Authority without observation is not governance — it is a blindfolded controller applying inputs to a system it cannot see, sometimes by chance stabilising it and sometimes accelerating its collapse.

The Ostrom jump is a variety jump

Ostrom's eight design principles for enduring commons institutions have been extensively studied and broadly validated. The principles include 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.

What the simulation adds to this literature is a formal characterisation of why these principles produce better outcomes. The performance gap between Architecture C (market) and Architecture D (community commons) is not attributable to stronger values, more committed actors, or better enforcement — all of these are held constant in the simulation. It is attributable to the shift from single-dimension to multi-dimension observation.

Ostrom's design principles are, functionally, a set of rules for opening and maintaining multiple observation channels simultaneously. Clearly defined boundaries determine who participates in the observation network. Monitoring creates the formal observation process. Collective-choice arrangements aggregate the distributed observations of community members into governance decisions. Graduated sanctions create feedback loops that maintain the integrity of the observation system itself — users who over-extract are sanctioned, which preserves the trust that makes monitoring effective.

The spatial distribution signal — one of the dimensions that distinguishes Architecture D from C — is only accessible because the community governance system maps extraction across its spatial territory, not just in aggregate. Individual community members, fishing or harvesting in different locations, observe the local state of different patches. The community governance system aggregates these distributed observations into a spatial picture that a central authority observing aggregate statistics simply cannot form. This is distributed sensing — the community as a sensor network.

The social pressure signal — the third dimension in Architecture D — is a governance-specific observation channel with no analogue in market or state management systems. When community members observe each other's behaviour, the social signal carries information about norm compliance, extraction pressure, and emerging conflicts that would not appear in any biological or economic indicator. This signal is available only because governance is embedded in a social system with ongoing relationships and mutual accountability.

The slow variable jump and intergenerational knowledge

The performance gap between Architecture D and Architecture E is smaller in absolute magnitude than the gap between C and D, but it is more significant in kind. The difference between 30% and 3.6% collapse risk is substantial, but more important is the qualitative difference in what each architecture can see.

Architecture D, despite its superior local monitoring, is effectively blind to the slow decadal trend in carrying capacity. Its observation window — monthly, with no formal mechanism for long-run baseline comparison — treats each year's conditions as independent of previous years. Slow trends are indistinguishable from normal variability at the timescale of any single community member's tenure in the governance system.

Architecture E avoids this blindness through access to the slow ecological signal dimension — operationalised in the simulation as a direct observation of the long-run carrying capacity trend, available because the governance system has multi-decadal baseline data. In a real governance system, this signal dimension is provided by intergenerational knowledge: the accumulated observations of previous generations encoded in cultural memory, governance protocols, oral tradition, and land-use practice.

The slow variable tracking panel in Figure 1 makes this visible. Only Architecture E's 24-month rolling mean shows any correspondence to the true carrying capacity trajectory. All other architectures discover the slow decline only after their stock has collapsed — their rolling means fall along with the stock, not in advance of it. They are not detecting a trend; they are experiencing a crisis.

This is the formal basis for what ecologists call traditional ecological knowledge: not a cultural artefact or a political recognition, but a multi-generational observation record of slow-moving ecological variables that no short-run monitoring programme can replicate. The governance protocols of communities with centuries of continuous resource management encode observations about slow variables — decadal drought cycles, long-run species dynamics, threshold behaviour at extreme stock depletion — that modern monitoring programmes have simply not existed long enough to observe.

The implication for conservation management is direct: the slow variable dimension of commons governance cannot be reconstructed by improving the observation 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.

Equity as a structural co-product

The simulation's equity finding deserves emphasis because it contradicts a common assumption in commons governance theory: that sustainability and equity are in tension, and that governance systems must choose between them.

Architecture E — the highest-performing architecture on resource conservation — also has the lowest extraction inequality (Gini 0.032). Architecture C — the market mechanism — has the highest inequality (0.096) and near-complete resource collapse. Open access (A) appears equitable (Gini 0.018) but this reflects universal impoverishment rather than equitable distribution — all users are equally unable to extract from a collapsed resource.

The co-production of equity and sustainability in Architecture E is not coincidental. It follows from the same architectural properties that produce conservation performance. Multi-dimensional monitoring that observes spatial distribution across users makes unequal extraction visible. Strong social accountability creates pressure toward equitable practice. Governance rules matched to seasonal dynamics apply uniformly rather than advantaging well-capitalised users who can extract at any time. And the long-run baseline awareness that prevents slow-trend blindness also prevents the gradual accumulation of extraction advantages by powerful actors that often precedes commons collapse in real systems.

The governance system with the highest requisite variety is also the most equitable. This is not a coincidence — it is a structural property. High variety observation sees inequality as well as ecological stress. Low variety observation is blind to both.

Part V: Limitations

Logistic growth is a simplification

The resource dynamics in the simulation use logistic growth — a smooth, continuous model that captures the basic regeneration-depletion structure of renewable resources but omits several properties of real ecosystems that may be significant for governance conclusions.

Real ecosystems exhibit threshold effects and hysteresis: stock levels that fall below critical thresholds may shift the system to a qualitatively different regime from which recovery is not simply the reverse of decline. A fishery that collapses past a recruitment failure threshold may not recover to previous productivity even if extraction is halted entirely. Logistic growth assumes continuous, reversible dynamics. In real systems, the consequences of brief excursions into the collapse zone may be far more severe than the simulation models. This suggests that the collapse risk figures in the simulation likely understate the real consequences of high-observation-latency governance.

Spatial dynamics in the simulation are simplified to nearest-neighbour diffusion at a fixed rate. Real resource systems have more complex spatial dynamics — directional flows, habitat-specific productivity, source-sink dynamics — that interact with governance architectures in ways the simulation does not capture.

Observation dimensionality is modelled as a scalar

The simulation assigns a single integer (observation dimensionality) to each architecture and uses this to parameterise the noise and accuracy of the observation. This is a useful simplification but misrepresents the actual structure of observation quality.

In reality, different observation dimensions have different properties: some are easily measured, some are difficult; some are available to all observers, some are available only to specifically positioned ones; some are reliable, some are noisy; some are leading indicators, some are lagging. The simulation treats all dimensions as equally informative and equally accessible to the governance system that nominally has them.

The slow ecological signal in Architecture E is particularly simplified. The simulation gives Architecture E direct access to the carrying capacity trend. Real traditional ecological knowledge is not a clean signal — it is embedded in cultural practices, seasonal protocols, and interpretive frameworks that require expertise to read correctly. The effective slow-variable observation that indigenous governance systems achieve is harder to access and easier to lose than the simulation's clean parametric assignment suggests.

Population growth and extraction pressure are held constant

The simulation holds the number of user groups and their baseline extraction needs constant over the thirty-year horizon. Real commons governance faces the additional challenge of governing under growing extraction pressure — more users, higher per-capita consumption, more powerful extraction technology. All governance architectures would perform worse under rising pressure; the question is whether the performance ranking would change.

The theoretical prediction — higher requisite variety architectures are more robust to pressure increases because they can detect and respond to rising pressure signals — is plausible but not demonstrated in this simulation. This is a significant omission given that population growth and technological intensification are the primary historical drivers of commons collapse.

Indigenous governance systems are diverse

Architecture E is a model of structural properties common to indigenous governance systems — long-run ecological embeddedness, high-variety observation, intergenerational knowledge, strong social accountability — not a model of any specific system. Real indigenous governance systems are enormously diverse in their specific protocols, decision rules, enforcement mechanisms, and ecological contexts.

The simulation demonstrates that the structural properties Architecture E embodies produce superior commons outcomes. It does not demonstrate that all indigenous governance systems have these properties to an equal degree, that no indigenous systems have failed historically, or that any specific indigenous system would perform identically to the model. Indigenous governance systems, like all governance systems, vary in quality, have been subject to historical disruption, and face challenges of adaptation to changed conditions.

The paper's claim is structural: governance systems with the requisite variety properties modelled in Architecture E will outperform those without them. It is an observation about architecture, not a blanket endorsement of all practices of all communities.

The epistemic dimension is not fully captured

Perhaps the most significant limitation is the simulation's implicit assumption that observation dimensionality is purely a matter of signal availability — that if a governance system has access to a signal dimension, it can use it effectively. This understates the epistemic complexity of traditional ecological knowledge.

Indigenous ecological knowledge is not simply a larger number of observations. It involves different epistemological frameworks for what counts as a relevant signal, how signals are interpreted, what relationships between signals are meaningful, and how observations translate into governance decisions.

These frameworks have been refined over generations precisely because they have been tested against ecological outcomes — the ones that led to collapse were abandoned; the ones that sustained resources were maintained and transmitted.

Replicating the observation dimensionality of traditional ecological knowledge without the interpretive framework that makes it legible is not obviously possible. An external monitoring programme that installed the same sensor array and collected the same data dimensions would not necessarily produce the same governance quality, because the signal processing — the knowledge of what the signals mean and how to respond to them — is embedded in the community's interpretive tradition, not in the signals themselves.

This limitation suggests that the simulation may overstate the achievable performance of any governance system that attempts to replicate the observation dimensionality of embedded indigenous governance through technological means alone.

Part VI: Implications

The four papers together

The four papers in this series have established four connected results using the same formal framework applied to different structural problems.

Paper one demonstrated the subsidiarity principle: localized disturbances cannot be stabilized by centralized controllers because aggregation destroys the spatial information needed for targeted response.

Paper two demonstrated the fractality principle: multi-scale disturbance environments cannot be stabilized by single-scale controllers because no single latency can cover all disturbance frequencies.

Paper three demonstrated the observability-democracy connection: citizen preferences cannot be reliably transmitted through deep representation chains because aggregation loss and noise accumulation destroy the signal before it reaches the policy layer.

This paper demonstrates the requisite variety principle: commons resource systems cannot be reliably governed by low-dimensionality observation systems because the resource's disturbance variety exceeds the governance system's observational capacity. The unobserved variety appears as uncontrolled variance — collapse events that the governance system cannot detect in advance or prevent.

The four findings share a common structure: aggregation destroys information, and destroyed information cannot be reconstructed downstream. Whether the information is spatial (paper one), temporal-frequency (paper two), preference-distributional (paper three), or ecological-multidimensional (this paper), the loss is irreversible and the governance consequences are structural. Institutional quality — better enforcement, better deliberation, better survey methodology, better compliance — operates on the signal after it arrives. It cannot recover the signal before it was destroyed.

What this means for resource governance policy

Current international frameworks for commons governance — the Convention on Biological Diversity, REDD+, fisheries agreements, water treaties — predominantly assign governance authority to state actors operating through administrative systems with the properties of Architecture B: high observation latency, single-dimension aggregate signals, and enforcement mechanisms calibrated to last year's data.

The simulation result — Architecture B performs worse than open access — is a direct challenge to this assignment of authority. It is not a challenge to governance per se. It is a challenge to the assumption that formal authority backed by aggregate monitoring constitutes adequate governance of complex renewable resources.

The implication is not that state governance should be withdrawn. It is that state governance needs to be redesigned around the requisite variety constraint. For resources with multi-scale disturbance environments — which is most of them — effective governance requires observation systems with dimensionality matching the disturbance variety. State governance systems can achieve higher observation dimensionality through co-management arrangements that integrate community monitoring into formal governance frameworks. The alternative — maintaining centralized authority with low-dimensionality observation — is the architecture the simulation shows is reliably worse than the baseline it was designed to improve upon.

Resource 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 are legitimate and important framings.

The simulation adds a third framing that does not depend on historical or cultural arguments: resource sovereignty is an engineering requirement for effective commons governance.

Communities with long-run embeddedness in a resource system have accumulated, through continuous observation across generations, the observation dimensionality required to govern that system across all relevant disturbance timescales. Displacing this governance with external administration that has Architecture B properties 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 governance protocols, seasonal practices, and land-use traditions is not transferred when authority changes hands. It is lost.

The reconstruction of this knowledge by external monitoring programmes is slow, expensive, and — as the limitations section notes — epistemically incomplete. 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 replaced by any amount of well-funded monitoring infrastructure operating on administrative timescales.

This makes the argument for indigenous resource sovereignty partly consequentialist rather than purely rights-based: communities embedded in their resource systems are observing and governing dimensions of those systems that no external authority can effectively monitor. Removing that governance destroys a public good — effective commons management — not merely a private right.

The framing should be precise: 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 — and that 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.

The connection to the GGF's bioregional architecture

The Global Governance Frameworks project structures sub-planetary governance around Bioregional Autonomous Zones — governance units defined by ecological boundaries rather than administrative ones, with authority over resource management within their territory. The requisite variety analysis provides formal grounding for this design choice.

Bioregional governance units, defined by watershed, ecosystem, or species range, correspond to the spatial extent over which resource dynamics are coupled. A governance system whose territorial boundary matches the resource system's spatial dynamics can observe the full resource state — including spatial distribution, source-sink dynamics, and boundary flows — that a governance system whose boundaries cross-cut the ecosystem cannot. Misaligned governance boundaries impose an artificial aggregation that destroys spatial information, reproducing the same information-loss mechanism at a different scale.

The GGF's principle that higher governance layers coordinate only what the local layer structurally cannot handle maps directly onto the requisite variety result. Local communities govern the dimensions of the resource system their proximity makes observable. Higher layers govern genuinely cross-boundary dynamics — migratory species, transboundary water flows, climate-driven regime shifts — whose observation requires governance systems operating at the appropriate spatial and temporal scale. The division is not political preference; it is the assignment of governance authority to the level that has the requisite variety to exercise it.

The commons, the climate, and the slow variable problem

Climate change is the commons problem at planetary scale: a shared atmospheric resource degraded by the aggregate of distributed extraction decisions, with governance structures that lack the observation dimensionality and temporal resolution to manage it effectively.

The slow variable problem the simulation identifies in commons governance has a direct analogue in climate governance. The slow 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 Architecture A through D in the simulation, climate governance discovered the slow variable in retrospect, as a crisis rather than as an early signal.

The communities most likely to have 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 communities' observations is not sentimental. It is the slow variable signal dimension that the planetary commons governance system most urgently needs and structurally lacks.

Part VII: Conclusion

The tragedy of the commons has a precise architectural cause: individual extraction decisions made without feedback from the collective resource state. Closing the feedback loop requires observation. The quality of the closure depends on the observation channel. And the observation channel is constrained by Ashby's Law of Requisite Variety: a governance system cannot stabilize a resource whose variety exceeds the governance system's observational capacity.

The simulation demonstrates this constraint quantitatively across five architectures facing identical resources and disturbances. The result is monotonic and structural: observation dimensionality determines governance performance. The finding that state management produces worse outcomes than open access is not a criticism of governance institutions — it is the predicted consequence of high observation latency and single-dimension aggregation applied to a multi-scale disturbance environment. It holds regardless of institutional quality. It holds regardless of how honest, diligent, or well-resourced the state management system is.

The Ostrom jump — from single-dimension to multi-dimension observation — reduces collapse risk from 86% to 30%. Ostrom's design principles are not better values. They are a set of institutional arrangements that open multiple observation channels simultaneously, increasing the governance system's effective variety to cover the fast and medium disturbance bands that single-dimension observation misses.

The slow variable jump — adding access to decadal ecological signal dimensions — reduces collapse risk from 30% to 3.6%. This is the dimension that traditional ecological knowledge provides and that no monitoring programme operating on administrative timescales can replicate. A governance system without access to the slow variable signal will discover slow trends retrospectively, as crises, rather than prospectively, as signals requiring early response.

These results connect this paper to the preceding three in the series. Paper one showed that centralized controllers destroy spatial information. Paper two showed that single-scale controllers cannot cover multi-frequency disturbance environments. Paper three showed that deep representation chains destroy preference signals. This paper shows that low-dimensionality observation systems cannot govern high-variety resource systems. The same mechanism — information loss in observation channels — produces governance failure in each domain.

The implication for commons governance policy is that observation quality is the prerequisite, not the add-on. Authority, enforcement, and institutional quality matter — but they matter on the signal that arrives at the governance layer. If the observation channel is too slow, too narrow, or too aggregated to distinguish the states the resource system occupies, improved institutional quality operates on noise rather than signal.

The implication for resource sovereignty is structural: communities with long-run ecological embeddedness have accumulated, through continuous intergenerational observation, the observation dimensionality required to govern the resource systems they inhabit across all relevant disturbance timescales. This is not a cultural claim or a historical claim. It is an architectural observation with the same formal status as the other findings in this series. Governance systems positioned within their resource systems, observing them continuously across multiple dimensions and across generations, have the requisite variety that externally positioned governance systems structurally lack.

The engineering diagnosis does not resolve the political questions — who holds authority, what rights are recognised, how governance transitions are managed. But it sharpens the terms of the question. The choice between governance architectures is not only a choice between institutional values or historical claims. It is a choice between observation systems — between governance that can see the resource and governance that cannot. The simulation is unambiguous about which produces better outcomes.

Appendix A: Mathematical formulations

Resource dynamics

Each of the $N_PATCHES = 12$ resource patches evolves according to discrete-time logistic growth with diffusion and extraction:

$$\begin{aligned}
 R_p(t+1) = & R_p(t) + r \cdot R_p(t) \cdot (1 - R_p(t)/K(t)) \\
 & + \beta \cdot \sum_{q \in N(p)} (R_q(t) - R_p(t)) \\
 & + \varepsilon_p(t) \\
 & - E_p(t)
 \end{aligned}$$

Where:

- $r = 0.08$ is the intrinsic growth rate
- $K(t)$ is the time-varying carrying capacity
- $\beta = 0.02$ is the diffusion coefficient between neighbouring patches
- $N(p)$ is the set of patches adjacent to patch p ($|i-j| \leq 2$)
- $\varepsilon_p(t) \sim N(0, \sigma_{fast}^2)$ is the fast stochastic shock
- $E_p(t)$ is extraction allocated to patch p

Stock is clipped to $(0, 1.5 \cdot K(t))$ at each step.

Multi-scale carrying capacity

The carrying capacity varies across three simultaneous disturbance bands:

$$K(t) = K_{base} + A_{med} \cdot \sin(2\pi \cdot t/P_{med}) + A_{slow} \cdot \sin(2\pi \cdot t/P_{slow})$$

Where:

- $K_{base} = 100$ (baseline per patch)
- $A_{med} = 8, P_{med} = 12$ months (seasonal cycle)
- $A_{slow} = -20, P_{slow} = 240$ months (decadal decline)
- $\sigma_{fast} = 3.0$ (fast stochastic shock per patch per step)

The slow component produces a net carrying capacity decline of approximately 20 units at peak (around month 120), representing long-run environmental degradation.

Ashby variety analysis

Let D denote the disturbance process driving the resource and R the regulator (governance system). The resource system's disturbance variety can be decomposed by frequency band:

$$V(D) = V(D_{\text{fast}}) + V(D_{\text{med}}) + V(D_{\text{slow}})$$

Where $V(\cdot)$ denotes \log_2 of the number of distinguishable states (Shannon variety). The governance system's variety is:

$$V(R) = \log_2(\text{obs_dims} \cdot \text{obs_resolution} / \text{obs_latency_penalty})$$

The requisite variety condition $V(R) \geq V(D)$ is met when the governance system can distinguish all states of the resource that require different governance responses. In practice, this requires:

- $\text{obs_latency} \leq 1/f_{\text{fast}}$ (latency below the fast disturbance period)
- $\text{obs_dims} \geq 2$ to distinguish seasonal from baseline states
- $\text{obs_dims} \geq 3$ with multi-decadal baseline to distinguish slow trends from variability

Only Architecture E satisfies all three conditions in the simulation.

Feedback loop integrity

Feedback loop integrity $\text{FLI} \in (0, 1)$ measures the degree to which extraction decisions are coupled to the current resource state:

$$\text{FLI} = \text{corr}(E(t), R_{\text{obs}}(t)) \cdot (1 / \text{obs_latency_penalty}) \cdot \text{obs_dims_factor}$$

Where:

- $\text{corr}(\cdot, \cdot)$ is the Pearson correlation between extraction and observed stock
- $\text{obs_latency_penalty} = 1 + \text{lag}/T$ (lag in months, $T = 360$)
- $\text{obs_dims_factor} = \min(\text{obs_dims} / V_{\text{disturbance_bands}}, 1.0)$

Open access (A) has $\text{corr} < 0$ (extraction rises with stock) but $\text{FLI} \approx 0.15$ because there is no aggregate coordination. State management (B) has structured quotas but $\text{FLI} \approx 0.12$ due to the latency penalty. Architecture E achieves $\text{FLI} \approx 0.78$ — the only architecture with substantial feedback loop integrity across

all disturbance bands.

Extraction allocation across patches

Total extraction by user group u at time t is allocated to patches proportional to current patch stock:

$$E_p(t) = E_{\text{total}}(t) \cdot R_p(t) / \sum_q R_q(t)$$

This represents users preferentially extracting from more productive patches — a realistic assumption for mobile extractors (fishing vessels, herders, foragers).

Gini coefficient

Extraction inequality at each time step is measured by the Gini coefficient across user groups:

$$G(t) = \sum_i \sum_j |E_i(t) - E_j(t)| / (2 \cdot N \cdot \sum_i E_i(t))$$

Mean Gini over the simulation (excluding 10-step warmup) provides the summary inequality metric reported in the results table.

Collapse risk

Collapse risk is the fraction of time steps in which total stock falls below the collapse threshold:

$$CR = (1/T) \cdot \sum_t \mathbb{1}[\sum_p R_p(t) < \theta \cdot K(t) \cdot N_{\text{patches}}]$$

Where $\theta = 0.20$ is the collapse threshold fraction and $\mathbb{1}(\cdot)$ is the indicator function.

Full simulation parameters

Parameter	Value	Notes
N_patches	12	Spatial resource patches
N_users	20	User groups
T	360	Time steps (months — 30 years)
K_base	100.0	Baseline carrying capacity per patch
r_growth	0.08	Intrinsic growth rate per step
β	0.02	Diffusion coefficient between patches
σ_{fast}	3.0	Fast stochastic shock std dev
A_med	8.0	Seasonal amplitude
P_med	12	Seasonal period (months)
A_slow	-20.0	Slow trend amplitude
P_slow	240	Slow trend period (months)
$\theta_{collapse}$	0.20	Collapse threshold (fraction of K)
Random seed	42	For reproducibility
Warmup	10	Steps excluded from metrics

Architecture observation parameters (obs_lag, obs_dims, quota_rigidity, sanctioning, slow_signal):

	A	B	C	D	E
obs_lag	0	12	3	1	1
obs_dims	1	1	1	3	6
quota_rigidity	0.0	0.7	0.0	0.9	0.95
sanctioning	0.0	0.3	0.0	0.8	0.9
slow_signal	False	False	False	False	True

Appendix B: Code and reproduction

Source code

The v6 simulator extends the series to the commons governance domain, replacing the preference-transmission model of v5 with a resource stock dynamics model governed by five distinct feedback architectures. It is implemented in Python using NumPy and Matplotlib.

The full source code is available at:

github.com/BjornKennethHolmstrom/ggf-governance-simulator (<https://github.com/BjornKennethHolmstrom/ggf-governance-simulator>)

The repository now includes six simulator versions:

File	Paper	Description
<code>ggf-simulator-v2.py</code>	Paper I	Single-node scalar feedback model
<code>ggf-simulator-v3.py</code>	Paper I	Ten-node vector model, localized shock
<code>ggf-simulator-v3-unadjusted.py</code>	Paper I	v3 with unstable gain — instability demo
<code>ggf-simulator-v4.py</code>	Paper II	Multi-scale disturbance, three architectures
<code>ggf-simulator-v5.py</code>	Paper III	Representation chain observability, four architectures
<code>ggf-simulator-v6.py</code>	Paper IV	Commons governance, requisite variety, five architectures

Reproducing the results

```
git clone https://github.com/BjornKennethHolmstrom/ggf-governance-simulator
cd ggf-governance-simulator
pip install numpy matplotlib
python ggf-simulator-v6.py
```

The simulation is seeded (`numpy.random.default_rng(seed=42)`). Default parameters exactly reproduce Figure 1 and the quantitative summary table in Part III.

Key architectural differences from v5

v6 is a structural departure from v5's preference-transmission model. Where v5 modelled information flow through a representation chain, v6 models a resource stock governed by five distinct feedback architectures with different observation and control properties.

Resource state space. Rather than citizen preference vectors, the state is a 12-patch resource stock evolving under logistic growth with diffusion, multi-scale disturbances, and extraction by 20 user groups.

Observation channel. The `observe_resource` function implements delayed, possibly aggregated, possibly noisy observation of the resource state. Architecture B observes with a 12-step lag; E observes with a 1-step lag and accesses the slow carrying capacity signal directly.

Extraction decisions. The `compute_extraction` function implements five qualitatively different extraction logics: unconstrained individual optimisation (A), quota compliance under annual survey (B), price-responsive extraction (C), community rule compliance with graduated sanctions (D), and seasonally-adaptive extraction with strong social accountability (E).

Multi-scale disturbances. The `carrying_capacity` function generates the full three-band disturbance environment. The slow trend is only observable to Architecture E, which has `slow_signal=True`.

Spatial dynamics. Patch connectivity is generated by `generate_patch_connectivity`, which creates a banded diffusion matrix. Extraction is allocated across patches proportional to current stock.

Modifying the parameters

Slow trend magnitude. Increasing `SLOW_AMP` (default 20) makes the decadal carrying capacity decline larger, which increases the performance gap between E and D. The gap between B and A also widens, because a larger slow trend means B's lagged quotas are further behind current conditions.

Observation lag (Architecture B). Reducing `obs_lag` from 12 to 6 substantially improves B's performance. This explores the question: at what observation frequency does state management become beneficial rather than harmful? The crossover point is architecturally informative.

Sanctioning strength. Reducing D and E's sanctioning parameters toward A and B's levels reveals how much of the community commons performance advantage is attributable to sanctioning versus observation dimensionality alone. The answer: most of the advantage persists even with weak sanctioning, confirming that observation dimensionality is the primary driver.

Collapse threshold. Increasing θ from 0.20 to 0.35 makes collapse easier to trigger and shifts all architectures' collapse risk upward. The relative ranking is preserved, but the absolute gaps change.

Number of users (N_users). Increasing N_users while holding resource parameters constant raises extraction pressure across all architectures. This explores the commons governance problem under population growth — a significant real-world extension not modelled in the base simulation.

Extending the model

Hysteresis and regime shifts. Replacing logistic growth with a bistable dynamics model (e.g. a double-well potential in stock space) would test whether the performance ranking is preserved when collapse has irreversible consequences. The prediction: the gap between E and all other architectures widens substantially, because the slow variable awareness that only E has is precisely what allows early intervention before threshold crossing.

Adaptive governance. Allowing governance parameters to evolve over time — architectures that learn from outcomes and adjust their observation systems — would model the emergence of adaptive management. The prediction: architectures with higher initial observation dimensionality adapt faster because they can observe what their governance decisions produce.

Multiple resource types. Extending to coupled resource systems (fishery + forest + water) with interactions between them would model biodiversity governance. The prediction: the observation dimensionality required for stable governance scales with the number of coupled resource types, further widening the advantage of embedded community governance over administrative systems with fixed observation channels.

Contributing

The repository is open source under MIT license. Extensions, empirical applications, and critiques are welcome via GitHub.

Appendix C: References and sources

A note on methodology

As with the preceding papers in this series, the concepts here were developed through extended dialogue with multiple AI systems — Claude (Anthropic), ChatGPT (OpenAI), Gemini (Google), DeepSeek, and Grok (xAI) — rather than through direct reading of the primary literature. The references below are the sources those systems identified as foundational, provided for readers who wish to engage with the primary literature directly.

The specific contribution of this paper — formalising commons governance as a requisite variety problem, demonstrating the state management paradox, and characterising traditional ecological knowledge as a slow-variable observation system — emerged from this collaborative process. The underlying theoretical tools belong to long-established traditions in cybernetics, ecology, and institutional economics. This paper brings them together with a simulation that makes the structural results quantitatively visible.

Cybernetics and requisite variety

Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman and Hall.

The Law of Requisite Variety is the central formal tool of this paper. Ashby's demonstration that only variety can absorb variety — and that a regulator cannot stabilize a system whose variety exceeds the regulator's own — provides the theoretical foundation for the observation dimensionality argument. Ashby discusses applications to biological and organizational systems; the application to commons governance is the contribution of this paper.

Ashby, W. R. (1952). *Design for a Brain*. Wiley.

The earlier treatment of adaptive systems and the mechanisms by which systems acquire requisite variety. Relevant to the discussion of intergenerational knowledge as an accumulated variety-acquisition process.

Wiener, N. (1948). *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT Press.

The foundational text of cybernetics. Wiener's treatment of feedback control in biological and social systems provides the conceptual background for the feedback loop integrity analysis.

Commons governance and institutional economics

Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.

The foundational empirical and theoretical work on self-governing commons institutions. Ostrom's eight design principles are reinterpreted in this paper as a set of observation channel opening mechanisms — institutional arrangements that increase effective observation dimensionality. The performance gap between Architecture C and D in the simulation corresponds to the observation dimensionality difference that Ostrom's principles create.

Ostrom, E. (2005). *Understanding Institutional Diversity*. Princeton University Press.

Extends the design principles to a broader institutional analysis framework (the Institutional Analysis and Development framework). Provides the theoretical grounding for understanding how governance rules interact with resource dynamics.

Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859), 1243–1248.

The original statement of the commons problem. This paper's central argument — that Hardin misidentified the cause as motivational rather than architectural — is a reinterpretation of Hardin's scenario in control-theoretic terms.

Dietz, T., Ostrom, E., & Stern, P. C. (2003). The struggle to govern the commons. *Science*, 302(5652), 1907–1912.

Overview of the conditions under which commons governance succeeds or fails. The observation quality conditions identified in this paper map onto several of Dietz et al.'s enabling conditions for successful governance.

Berkes, F. (1989). *Common Property Resources: Ecology and Community-Based Sustainable Development*. Belhaven Press.

Early systematic documentation of community-based natural resource management. Berkes' concept of adaptive co-management anticipates the observation dimensionality argument: effective co-management works because it integrates the observation systems of multiple positioned actors.

Traditional ecological knowledge

Berkes, F. (2008). *Sacred Ecology*. 2nd ed. Routledge.

The most comprehensive treatment of traditional ecological knowledge as a governance resource. Berkes' account of TEK as an adaptive management system operating across multiple timescales maps directly onto the slow-variable observation argument of this paper.

Gadgil, M., Berkes, F., & Folke, C. (1993). Indigenous and traditional knowledge of the environment. *Ambio*, 22(2–3), 151–156.

A foundational paper establishing traditional ecological knowledge as a form of ecological monitoring and governance that operates across timescales unavailable to modern science. The argument that TEK provides information about slow variables is directly relevant to the Architecture E findings.

Tengö, M., Brondizio, E. S., Elmqvist, T., Malmer, P., & Spierenburg, M. (2014). Connecting diverse knowledge systems for enhanced ecosystem governance. *Ambio*, 43(5), 579–591.

Proposes a "multiple evidence base" framework for integrating indigenous and local knowledge with scientific knowledge in biodiversity governance. This framework is consistent with the observation dimensionality argument: different knowledge systems are not competing narratives but complementary observation channels covering different signal dimensions.

Ens, E. J., et al. (2015). Indigenous biocultural knowledge in ecosystem science and management. *Biological Conservation*, 181, 133–149.

Empirical review of how indigenous biocultural knowledge contributes to biodiversity conservation outcomes. Documents the slow variable signal dimensions that long-run community observation provides.

Ecology and resource dynamics

May, R. M. (1977). Thresholds and breakpoints in ecosystems with a multiplicity of stable states. *Nature*, 269(5628), 471–477.

The foundational treatment of ecological threshold effects and regime shifts — the nonlinear dynamics that logistic growth simplifies away but that make the slow variable detection argument more urgent in real systems.

Folke, C., Carpenter, S., Walker, B., Scheffer, M., Elmqvist, T., Gunderson, L., & Holling, C. S. (2004). Regime shifts, resilience, and biodiversity in ecosystem management. *Annual Review of Ecology, Evolution, and Systematics*, 35, 557–581.

Comprehensive treatment of regime shifts — the irreversible threshold crossings that the logistic growth model understates. The argument that governance must detect slow variables before regime shift crossing is made more urgent by this literature.

Holling, C. S. (1978). *Adaptive Environmental Assessment and Management*. Wiley.

The foundational text of adaptive management — the framework that most closely anticipates the observation dimensionality argument in an ecological context. Holling's emphasis on the importance of monitoring across multiple spatial and temporal scales maps directly onto the requisite variety analysis.

Spatial governance and bioregionalism

Sale, K. (1985). *Dwellers in the Land: The Bioregional Vision*. Sierra Club Books.

The canonical statement of bioregionalism as a governance philosophy. Sale's argument that governance units should be defined by ecological rather than administrative boundaries is given a formal basis in this paper's analysis of observation dimensionality and spatial coupling.

McGinnis, M. D., & Ostrom, E. (2014). Social-ecological system framework: Initial changes and continuing challenges. *Ecology and Society*, 19(2), 30.

The Social-Ecological System (SES) framework as an integrative approach to analysing commons governance. The SES framework's emphasis on matching governance structures to resource system properties is consistent with the requisite variety argument.

Indigenous sovereignty and governance

Cobo, J. M. (1983). *Study of the Problem of Discrimination Against Indigenous Populations*. United Nations Economic and Social Council.

The foundational UN document establishing the international framework for indigenous rights. The paper's argument that resource sovereignty is an engineering requirement rather than solely a rights-based claim supplements rather than replaces the rights framework Cobo develops.

Anaya, S. J. (2004). *Indigenous Peoples in International Law*. 2nd ed. Oxford University Press.

The comprehensive legal treatment of indigenous peoples' rights in international law, including resource rights. Provides the legal context within which the engineering argument operates.

Tipa, G., & Welch, R. (2006). Cultivating a culture of participation: Maori involvement in local government decision-making. *Planning Theory and Practice*, 7(1), 61–78.

A case study of Maori participation in water governance in New Zealand — a practical example of the integration of traditional ecological knowledge into formal governance frameworks that the analysis of Architecture E models in stylised form.