

**Gluing Leviathan:  
A Local-to-Global Theory of  
Civilizational Reconstruction**

PIATRA . INSTITUTE

May 2026

## Abstract

A civilization is rarely planned and never built whole. It assembles from local rule systems, including a village's land tenure, a court's enforcement habits, a militia's threshold for force, a market's accounting conventions, and a monastery's grain reserve, whose interfaces with one another either glue into a coherent macro-order or fail to. This paper formalises the gluing condition. A post-collapse landscape is modelled as a finite cell complex whose vertices are settlements and whose higher cells are the routes, watersheds, debts, and jurisdictions that bind them. To each region we assign a finite space of admissible institutional packages drawn from seven dimensions: property, labor, money, contract enforcement, violence, resource governance, and knowledge retention. Restriction maps record what each local package commits to at every shared edge. A *global section* of the resulting cellular sheaf is a coherent macro-order. The first cohomology  $H^1$  of the sheaf measures the obstruction to the order existing globally. Above this compatibility layer sits a dynamical layer whose attractors are the historically recurrent macro-regimes: kin-commons, warlord tribute, neo-feudal land control, monastic redistribution, guild-merchant order, state-command, capitalist market order, and commons federation. Both layers are computed explicitly. A worked three-settlement river system shows that of three plausible repair operations, a water compact, an exchange-rate table, and a safe-conduct treaty, only the first two are necessary to admit any global section at all. The third sits on an edge that was never the obstruction. It shrinks the cochain spaces without buying compatibility, the kind of error that institutional reformers make routinely. A four-dimensional parameter sweep over violence, transport, energy, and literacy finds that the capitalist attractor wins 3.1% of the parameter cube, while the neo-feudal and state-command basins together win 43.3%. These shares are facts about the model's geometry, not estimates of historical frequency. Capitalism's stability is a geometric coincidence of low violence, secure transport, abundant energy, and preserved literacy. Reconstruction, on this view, is the engineering of interfaces under a parameter regime that may or may not permit the macro-order being aimed at.

## 1. What survivors rebuild

Civilizations have collapsed before, sometimes catastrophically and sometimes by long unwinding, and the historical record of what survivors rebuild is unevenly studied but real. Roman provincial cities became manors, monasteries, and bishoprics in the centuries after the Western Empire dissolved, and the trade fairs of Champagne were five centuries away when the last Roman aqueduct in Gaul ran dry. The Western Han bureaucratic state gave way to the Three Kingdoms warlords, who in turn gave way to a re-imposed imperial bureaucracy under the Sui and Tang. Mayan polities collapsed into village agriculture that lasted longer than the polities had. The post-Bronze-Age Mediterranean spent four hundred years as scattered iron-using villages before the Greek poleis assembled. Whatever post-collapse societies tend to do, they reach the new macro-order through intermediate forms that look little like either the order they collapsed from or the order a contemporary reader might expect them to head toward.

Three answers to *why* circulate widely. The first is that capitalism is the natural attractor of any rebuilding society because humans like to trade. The second is that the strongest warlord wins,

with everything else following from extraction. The third is that institutions are cultural choices and the survivors will rebuild whatever their inherited norms predispose them to. Each answer captures something. None survive contact with the variety in the historical record. Trade is one ingredient among many in a capitalist macro-order, alongside alienable property, enforceable contracts, mobile labor, credit memory, reinvestable capital, and secure routes; most of those have to be in place at once for the assembly to hold. Warlords win locally and lose at scale, because extraction without production hits a ceiling that armed men cannot push through. Cultural inheritance is real, and it sits inside ecological, energetic, and infrastructural constraints that determine which inherited norms can scale up and which collapse on first contact with the territory.

The frame this paper proposes shifts the level of analysis. Survivors solve local coordination problems. They decide who owns the river bank, who enforces a debt, who travels armed, what counts as money, who eats first when stores run low, who is permitted to read. A macro-order exists when those local decisions become *mutually compatible* across the territory and *dynamically stable* under environmental pressure. Compatibility is the question of whether a particular pattern of local rules can coexist at all without contradiction at the seams between regions. Stability is the question of whether a compatible pattern, once reached, persists under the shocks the environment delivers. The two questions are usually conflated and need to be separated. A macro-order can fail because compatibility fails, leaving the order impossible. It can also fail because dynamics is hostile, leaving the order merely fragile. The responses to those two failures differ.

For compatibility we use the language of cellular sheaves, an offspring of algebraic topology that has matured over the past fifteen years into a tool for reasoning about distributed agreement on networks (Curry, 2014; Robinson, 2014; Hansen & Ghrist, 2019; Edelsbrunner & Harer, 2010). For dynamics we use attractor landscapes in the older sense of Holling (1973) and Strogatz (2015), together with the more recent vocabulary of social-ecological resilience (Walker, Holling, Carpenter & Kinzig, 2004; Folke, 2006) and critical transitions (Scheffer, 2009). Neither tradition has been brought to bear on civilizational reconstruction in the form proposed here.

Institutional economics, from North (1990) through North, Wallis & Weingast (2009) and the medieval trade studies of Greif (2006), establishes that rules and enforcement matter, and supplies the historical case material that disciplines any modelling exercise. Williamson (2000) gives the analytic vocabulary of transaction costs and contract enforcement that grounds the labor- and contract-regime distinctions used here. Empirical comparative political economy, in particular Acemoglu, Johnson & Robinson's (2001) study of colonial-origin institutional persistence and the popular synthesis in Acemoglu & Robinson (2012), shows that institutional differences persist over centuries and predict economic divergence in ways that pure geography or resource-endowment accounts cannot. Tilly (1990) gives the parallel claim on the violence side: European state formation co-evolved with capital accumulation and coercive capacity through tightly coupled feedback loops that the present model treats as the dynamics layer above the sheaf. Turchin (2003) frames the longer trajectory through secular cycles in agrarian states, in which surplus accumulation, elite over-production, and demographic pressure drive recurrent collapse-and-reconstruction patterns over centuries. Commons governance, since Ostrom (1990), shows how durable non-state, non-

capitalist regimes manage overlap rules at the level of design principles.

The contribution here is to compose these traditions into a single object: a sheaf whose stalks carry institutional configurations, whose restriction maps carry overlap-agreement constraints, whose cohomology localises the obstruction, and whose set of consistent macro-orders is then filtered through an attractor landscape over the four parameters that the historical record on collapse most consistently identifies as load-bearing: violence, transport, energy, and literacy.

## 2. The construction

### 2.1 Cells, stalks, restrictions

*Notation.* Throughout the paper  $X$  denotes a finite cell complex;  $v, e, f$  index its vertices, edges, and faces;  $\mathcal{F}$  is the institutional sheaf;  $\rho_{UV}$  is the restriction map from a region  $U$  to a sub-region  $V$ ;  $C^k(X, \mathcal{F})$  is the  $k$ -th cochain space and  $H^k(X, \mathcal{F})$  is its  $k$ -th cohomology;  $s$  ranges over sections, local or global; the seven institutional dimensions  $P, L, M, C, V, R, K$  are typeset in upright capitals throughout, with their components italicised on first use. Greek letters carry their standard analytic-topological meaning ( $\rho$  for restriction maps,  $\beta$  for the inverse temperature in the Boltzmann sharpness of §3).

A post-collapse landscape is modelled as a finite cell complex  $X$  whose elements correspond to social geography, infrastructure, and jurisdiction, layered together. Vertices are the places where institutional decisions are made and enforced: settlements, forts, farms, monasteries, workshops, clinics, mines. Edges are the relations along which decisions in one place have to be reconciled with decisions in another: roads, rivers, trade routes, kinship ties, postal networks, radio links. Faces and higher cells are the larger bounded regions over which a single rule set can claim to apply: watersheds, legal jurisdictions, defended marketsheds, ecclesiastical provinces.

Modelling society this way treats institutional life as a graph of interfaces whose dimensions matter as much as its nodes. Two villages connected by a long but well-policed road are *closer*, in the sense that matters here, than two villages adjacent across an unpoliced ridge. A monastery that issues a unified set of grain-storage rules to a dozen villages occupies a face that those villages share, even when no two of them touch. The cell complex is the formal record of which regions have to agree on something with which other regions, and at what dimension of the social order.

The substrate is *given* in the model. Ecology, weapons technology, demography, climate, and the residue of pre-collapse infrastructure determine which vertices and edges a given landscape contains. The model takes those as parameters; it does not endogenise them. The restriction is deliberate. Institutions can sometimes alter their substrate by building new roads, draining new fields, or federating new defence networks, but on the time-scale of a single reconstruction episode the substrate is mostly inherited rather than chosen.

To each region  $U \subseteq X$  we assign a finite set of admissible institutional configurations:

$$\mathcal{F}(U) = P_U \times L_U \times M_U \times C_U \times V_U \times R_U \times K_U,$$

where the seven factors are the seven institutional dimensions catalogued in §2.2. A *local section*  $s_U \in \mathcal{F}(U)$  is one coherent institutional arrangement inside  $U$ . The seven-dimension product is the universe of possibilities; vertex profiles further restrict each region to the subset of configurations that ecology, history, and existing power structures actually admit there.

When  $V \subseteq U$  is a sub-region, typically an edge contained in two vertex regions, a restriction map

$$\rho_{UV} : \mathcal{F}(U) \rightarrow \mathcal{F}(V)$$

projects an institutional package onto the dimensions that matter on  $V$ . A road requires the regions it connects to agree on what kind of violence is licit between them and how their contracts are enforced; the restriction map sends a full institutional package to its (violence-rule, contract-rule) component. A river requires agreement on water rights and projects accordingly. The choice of restriction map for each edge is a modelling decision of substantive content, and it carries its own argument. A modeller who treats a road as carrying only a violence rule is claiming, implicitly, that other dimensions of social life are uncontested at the road's interface, a claim that may or may not be realistic about a particular historical road. The model exposes the choice and lets it be argued.

**Definition 1** (Global section). *A global section of  $\mathcal{F}$  is a family  $s = \{s_v \in \mathcal{F}(v)\}$  indexed by the vertices of  $X$  such that for every edge  $e$  with endpoints  $u, v$  the projections agree:*

$$\rho_{u,e}(s_u) = \rho_{v,e}(s_v).$$

A global section is a coherent macro-order: every settlement chooses a local institutional configuration that survives every interface check with every neighbour with which it shares an edge. The existence of even one global section is non-trivial. For a typical post-collapse complex with diverse local institutional histories, no global section exists at all without targeted institutional repair work. The historical record bears this out repeatedly. The Carolingian renaissance was, structurally, an attempt to install a uniform legal and ecclesiastical layer over local custom that had previously failed to glue at almost every overlap. The Hanseatic League was a merchant-court protocol that allowed cities with mutually incompatible municipal law to nevertheless reconcile on the trade interface.

## 2.2 Seven dimensions

The seven institutional dimensions correspond to coordination problems that historical reconstruction settings demonstrably have to solve, and the admissible values for each dimension correspond to recurring solutions of those problems in the historical record.

**Property regime** ( $P$ ) decides who has the right to exclude others from a thing or a place. The four admissible values, commons, clan ownership, private alienable property, and state ownership, are the historically dominant patterns. Commons (managed open-access regimes in the Ostrom (1990) sense) are stable over centuries when the design principles for collective monitoring are intact.

Clan ownership embeds property in kinship networks and is resilient when those networks are dense. Private alienable property requires courts that recognise transfers across non-kin and is therefore cognitively and institutionally expensive. State ownership requires a coercive apparatus capable of enforcing claims on behalf of the centre. The four are distinct attractors with different interface requirements rather than points on a single axis.

**Labor regime (*L*)** decides how productive activity is mobilised. Five values are used: household labor (the dominant form for most of human history), wage labor (which presupposes mobile labor markets and impersonal exchange), serfdom (which binds labor to land), slavery (which alienates labor entirely and treats persons as property), and cooperative labor (which pools labor without alienation). The labor regime interacts strongly with property and contract regimes; serfdom needs land to be alienable in some form, and slavery needs courts willing to recognise persons as property.

**Money or accounting (*M*)** decides how value is recorded across time and traded across distance. Four values: barter, grain ledgers (storage-backed accounting in the manner of late Bronze Age temple economies), metal coin, and credit. Greif (2006) and the medieval-trade literature it draws on demonstrate that credit predates the banking forms commonly associated with it; merchant communities in the eleventh and twelfth centuries ran credit networks long before legal infrastructure for impersonal lending existed.

**Contract enforcement (*C*)** decides how agreements are made to stick. Four values: kin-oath, guild membership, court adjudication, and armed enforcement. Each carries its own scale ceiling. Kin-oath does not extend beyond a few hundred people. Guilds extend to a profession or a trade circuit. Courts extend to a polity. Armed enforcement extends to whatever territory the armed party can reach.

**Violence regime (*V*)** decides who is permitted to use force, against whom, and under what threshold. Five values: raiding (uncoordinated predation), warlord (organised extraction with weak legitimacy), militia (community defence), monopoly (a state's claim to all licit force, in the Weberian sense), and federated defence (mutual-aid arrangements among polities). The North, Wallis & Weingast (2009) distinction between *limited-access orders* (which manage violence by restricting access to valuable organisations) and *open-access orders* (which manage violence through rule-bound competition among organisations) is precisely a structural claim about which combinations of values from this dimension can scale.

**Resource governance (*R*)** decides how non-substitutable scarce resources, including water, soil fertility, fuel, and mineral deposits, are allocated. Four values: common, private diversion, state allocation, and kin-allocated. Resource governance is where commons-style and capitalist-style macro-orders most commonly collide at their interfaces. An upstream private water diversion is incompatible with a downstream commons even when every other interface in the territory is glued.

**Knowledge retention (*K*)** decides how technical and administrative knowledge survives generations. Four values: oral tradition, scribal administration, technical manuals, and formal schools.

Knowledge retention is the slowest variable and often the decisive one. Tainter (1988) argues that complex societies collapse partly because the marginal productivity of additional administrative complexity falls; once it falls below the cost of maintaining the literacy infrastructure that complexity requires, the literacy infrastructure itself collapses, and what remains can support only simpler institutional regimes.

The Cartesian product has  $4 \cdot 5 \cdot 4 \cdot 4 \cdot 5 \cdot 4 \cdot 4 = 25,600$  candidate configurations. The size is small enough to enumerate exhaustively and large enough that combinatorial restrictions on its members are non-trivial. Vertex profiles in the worked example of §4 admit between 16 and 64 configurations per region; profiles in larger applications can be tighter or looser depending on how aggressively the modeller filters by historical and ecological constraints.

### 2.3 The obstruction class

The natural language for *whether and how local data assembles* is sheaf cohomology. The construction is concrete enough to compute on. Lift each finite stalk  $\mathcal{F}(c)$  to its real-coefficient indicator space  $\mathbb{R}^{|\mathcal{F}(c)|}$ , taking each admissible configuration as a basis vector. The restriction maps become  $\{0, 1\}$ -valued matrices: each row of  $\rho_{v,e}$  is the indicator of one element of the edge stalk  $\mathcal{F}(e)$ , and the entry in column  $j$  is 1 when the  $j$ -th element of  $\mathcal{F}(v)$  projects to that edge element. The cochain complex is

$$0 \rightarrow C^0(X, \mathcal{F}) \xrightarrow{d} C^1(X, \mathcal{F}) \rightarrow 0$$

with  $C^0 = \bigoplus_v \mathbb{R}^{|\mathcal{F}(v)|}$  and  $C^1 = \bigoplus_e \mathbb{R}^{|\mathcal{F}(e)|}$ . The coboundary  $d$  acts on each edge by the difference of its endpoint restrictions:  $d_e(c_u, c_v) = \rho_{v,e} c_v - \rho_{u,e} c_u$ .

Coefficient choice carries content. We compute over  $\mathbb{R}$  because the rank computations reduce to standard linear algebra. The construction has a parallel over  $\mathbb{Z}$ , where torsion in  $H^1(X, \mathcal{F}; \mathbb{Z})$  encodes obstructions invisible to the real-coefficient version: an obstruction class can vanish in  $\mathbb{R}$ -coefficients while a non-trivial  $\mathbb{Z}$ -torsion element prevents an integer global section. For institutional configurations the integer count of admissible global sections is the practically relevant question, and we therefore report it alongside the real-coefficient  $\dim H^1$  in the worked example. Edelsbrunner & Harer (2010) develop the broader homological framework on which both formulations rest.

**Definition 2** (Institutional obstruction class). *The institutional obstruction class is*

$$o(X, \mathcal{F}) \in H^1(X, \mathcal{F}) = C^1 / \text{im } d.$$

*Its dimension counts independent edge-disagreement directions that no choice of per-vertex configuration can erase. Its support across the edges of  $X$  identifies which interfaces are blocking integration.*

The interpretive content is direct. When  $\dim H^1 = 0$ , every overlap can be reconciled by some assignment of local configurations to vertices. When  $\dim H^1 > 0$ , at least one overlap struc-

turally resists every assignment, and the obstruction can be localised by examining which edges' contributions to  $C^1$  lie outside  $\text{im } d$ . The dimension of  $H^0 = \ker d$  counts the linearly independent consistent assignments in the lifted space; it is bounded below by the number of admissible global sections in the discrete sense, with the two numbers differing in general. A model can carry  $\dim H^1 > 0$  and still admit many integer global sections through redundancy across overlaps; a model can have  $\dim H^1$  drop after a repair without yet admitting any integer global section, which simply means that the repair has reduced the obstruction without yet eliminating it. Both diagnostics matter, and the worked example below reports both.

The applied-sheaf literature provides a growing toolkit for this kind of computation. Robinson (2014) develops sheaves over networks with explicit consistency-radius diagnostics. Curry (2014) gives the modern theory of cellular sheaves and cosheaves. Hansen & Ghrist (2019) develop a spectral theory that opens cellular sheaves to Laplacian and diffusion analysis, with direct uses in the dynamical layer of the present model that are left for future work. Hansen & Gebhart (2020) and Bodnar et al. (2022) extend the apparatus into machine-learning contexts under the names *sheaf neural networks* and *neural sheaf diffusion*, using cellular sheaf Laplacians to diffuse information over heterogeneous network domains. The institutional sheaf used here is closer in shape to the discrete-data settings of Robinson and Curry than to the differentiable settings of these later papers, and the computational toolkit transfers directly.

The construction has cousins in three nearby fields. Classical constraint-satisfaction (CSP) frameworks pose the same agreement-on-overlaps question in Boolean form, asking only whether a satisfying assignment exists. The cohomological reading refines that question: not only whether satisfaction is possible but which edges obstruct it, with  $\dim H^1$  counting independent obstruction directions and the support of the obstruction class localising them. Quantum contextuality in the sense of Abramsky & Brandenburger (2011) uses identical sheaf-cohomological apparatus on a different stalk content (probability distributions over measurement outcomes); the institutional construction here borrows the formal machinery while substituting institutional configurations for measurement outcomes. Distributed consensus and Byzantine agreement protocols (Riess & Ghrist, 2022) study a structurally similar local-to-global question for evolving network states; their sheaf-Laplacian gossip dynamics provide computational tools that the dynamical layer of the present model is positioned to inherit.

**Lemma 1** (Minimum repair localisation). *Let  $X$  be a cell complex with sheaf  $\mathcal{F}$ , and let  $E_b \subseteq E$  denote the set of edges  $e = (u, v)$  for which the restricted stalks have empty intersection on  $e$ :*

$$E_b = \{e = (u, v) \in E : \rho_{u,e}(\mathcal{F}(u)) \cap \rho_{v,e}(\mathcal{F}(v)) = \emptyset\}.$$

*A set of repair operations  $R$  admits a global section only if for every  $e \in E_b$  at least one operation  $r \in R$  widens an endpoint stalk on  $e$  such that the post-repair intersection on  $e$  is non-empty.*

The proof is contrapositive. If some  $e \in E_b$  remains untouched by any repair in  $R$ , then  $\rho_{u,e}(\mathcal{F}(u)) \cap \rho_{v,e}(\mathcal{F}(v)) = \emptyset$  in the post-repair sheaf as well. No assignment  $s_u \in \mathcal{F}(u)$ ,  $s_v \in \mathcal{F}(v)$  can satisfy  $\rho_{u,e}(s_u) = \rho_{v,e}(s_v)$ , because the required common projection does not

exist. The single edge  $e$  blocks every candidate global assignment, regardless of how compatibility is restored on every other edge. The Lemma converts the obstruction class from a diagnostic into a guide: a candidate repair that does not touch any  $e \in E_b$  cannot reduce the integer global-section count from zero. The worked example of §4 confirms this numerically by exhibiting an institutional repair (the safe-conduct treaty) that targets an edge already in  $E \setminus E_b$  and fails to enlarge the global-section count.

### 3. Eight attractors and what selects among them

The sheaf decides which macro-orders are *consistent*. Stability is a separate question, settled by a dynamical layer above the sheaf. The historical record on post-collapse societies, read as a record of dynamics, is a record of recurrent attractors: a small number of macro-regimes that turn up again and again across time and geography, with intermediate states unstable enough that societies pass through them quickly and rest in the recurrent ones for centuries. Turchin (2003) describes one component of the dynamics, secular cycles in agrarian states, in which surplus accumulation, elite over-production, and demographic pressure drive recurrent oscillation between expansion, stagflation, crisis, and depression. The present model treats such cycles as one mechanism by which a society moves around in attractor space rather than fixing on a single attractor permanently.

Eight attractors are used in this paper. They are stylised. An exhaustive taxonomy would distinguish more, and the eight cover the historically dominant patterns well enough for the basin geometry to be meaningful.

The **kin-commons** attractor is what subsistence settlements relax into when there is no surplus to fight over and no predation to fight off: small scale, household labor, kin oaths, common resource governance, oral knowledge. It dominates wherever human populations have had to recover from acute collapse with no exogenous reinforcement. The post-classic Maya village pattern is a clear case (McAnany & Yoffee, 2010, who also caution that the popular collapse-as-civilizational-failure narrative associated with Diamond (2005) flattens what was, on the ground, a long continuity of village agriculture). The post-Bronze-Age eastern Mediterranean villages traced by Cline (2014) are another.

The **warlord tribute** attractor is what insecure surplus settlements fall into when predation is high and nothing prevents organised violence from extracting from organised production. The early post-Roman north European countryside (Heather, 2006), the warlord period of post-Han China (Lewis, 2009), and the immediate aftermath of many twentieth-century state collapses are warlord-tribute regimes by this taxonomy.

The **neo-feudal** attractor is what warlord tribute hardens into once the extractor settles into a particular landscape: land control plus protection-for-rent, with serfdom or quasi-serfdom as the labor regime. Carolingian and post-Carolingian Western Europe is the canonical case (Wickham, 2005); Tokugawa Japan is another, with different surface forms and the same underlying gluing pattern.

The **monastic / temple redistribution** attractor appears when literacy survives in an institution that also controls grain storage and extends sacred legitimacy. Egyptian temple economies, the

early medieval Western monastic networks studied by Wickham (2005), and the Tibetan monastic estates are examples. Knowledge retention is the load-bearing dimension. Remove it and the regime collapses.

The **guild-merchant** attractor is what trading cities fall into when transport security and standard weights are high enough to support craft specialisation but no central state has consolidated the territory: trade law, weights and measures, court enforcement at the city level, federated defence among cities. The Hanseatic League and the late medieval Italian and German trading cities are the cleanest examples.

The **state-command** attractor appears when bureaucratic knowledge survives but private enforcement and trade networks do not. The Han, Tang, and Ming bureaucracies represent the durable form (Lewis, 2009, on the post-Han state reformation under Sui and Tang); the modern European nation-state, in Tilly's (1990) co-evolutionary account of coercion and capital, is a different but structurally analogous form; the modern Soviet bloc is a more recent and less durable example. The hallmark is a monopoly on legitimate violence combined with state allocation of resources and money.

The **capitalist market order** appears when alienable property, mobile wage labor, credit memory, court enforcement, monopolised but ordinary violence, secure transport, and reinvestable surplus are all simultaneously available. Early modern Northwest Europe is the case it is most often inferred from, and the inference is fragile: each of the conditions had to be in place at once, and the historical case material on each condition is contested. Acemoglu, Johnson & Robinson (2001) trace the divergent persistence of these conditions through the colonial-origins channel, in a way that the present model reads as differential basin reachability across the parameter cube.

The **commons federation** attractor appears when local institutions are strong, predation is low, and federation across local commons is mediated by mutual recognition rather than central enforcement. Late medieval Switzerland and the Iroquois Confederacy are partial examples; the design principles in Ostrom (1990) function as the gluing rules for commons federations.

The macro-state at time  $t$  in the model is the tuple

$$S_t = (X_t, \mathcal{F}_t, s_t, \theta_t)$$

where  $\theta_t$  contains the environmental parameters and the dynamics

$$S_{t+1} = T(S_t, \epsilon_t)$$

are driven by shocks  $\epsilon_t$  (famines, epidemics, invasions, charismatic founders, technical rediscoveries). The path-dependence and lock-in dynamics that complexity economics has documented in market evolution (Arthur, 1994) operate at the institutional level too: an attractor reached early constrains the attractors reachable later. We do not specify  $T$  at the level of agent dynamics in this paper. We specify, instead, a transparent fitness function for each attractor over the four-dimensional

parameter cube  $[0, 1]^4$  of (violence intensity, transport security, energy surplus, literacy retention), and we read the *basin of attraction* at a parameter point as the share of a Boltzmann distribution over the eight fitnesses with inverse temperature  $\beta = 12$ . The dominant attractor at a point is the argmax of those shares.


The choice of  $\beta = 12$  sets the sharpness of the soft argmax; it does not affect which attractor dominates a given point. A basin volume counts the grid cells on which an attractor is the fitness-maximiser, and the maximiser is independent of  $\beta$  for any  $\beta > 0$ , so the basin volumes reported below are exactly  $\beta$ -invariant: recomputing the sweep at  $\beta = 5$ ,  $\beta = 12$ , and  $\beta = 20$  returns identical volumes (capitalist 0.031, neo-feudal 0.249, state-command 0.184 in every case). What  $\beta$  governs is the mass that leaks onto the non-dominant attractors at each point, and hence the soft shares quoted at the corners below. As  $\beta \rightarrow \infty$  the distribution collapses to hard argmax (each grid cell goes one hundred percent to its winning attractor); as  $\beta \rightarrow 0$  it flattens to uniform  $1/8 = 0.125$  shares. The middle range used here reads as “the strongest attractor wins most of the basin mass at each point, with measurable but small mass leaking onto neighbours”. The treatment is deliberately legible: every parameter point produces a probability vector that can be read directly off the fitness functions, and the basin volumes reported in §5 reflect the geometry of those eight bumps in the parameter cube rather than any hidden stochastic dynamics. A richer extension would compose strategic interactions in the manner of compositional game theory (Ghani, Hedges, Winschel & Zahn, 2018), letting open games on the cells of  $X$  supply the dynamics directly; that extension is left for future work.

#### 4. Three settlements on a river

The smallest cell complex on which the framework does substantive work has three vertices and three edges. We use three settlements in distinct ecological and institutional positions:

- *A*, an upstream fortified settlement: private alienable land, wage and serf labor, metal coin, court and armed enforcement, warlord and militia violence, private upstream water diversion, scribal and manual knowledge.
- *B*, a downstream farming village: clan and commons land, household and cooperative labor, grain ledgers and barter, kin-oath and guild enforcement, militia and federated defence, common and kin-allocated water, oral knowledge.
- *C*, a trading town: private alienable land, wage labor, metal coin and credit, court and guild enforcement, monopoly and militia violence, private and state-allocated resource rules, scribal and manual knowledge.

The three edges carry distinct overlap requirements. The river *AB* requires agreement on the resource-governance dimension. The grain route *BC* requires agreement on the money dimension. The military road *AC* requires agreement on both the violence dimension and the contract-enforcement dimension. Each edge could in principle carry agreement on more dimensions; the choices here reflect what the historical literature suggests is the *load-bearing* contestation at each kind of interface.



Three repair operations are available. A *water compact* widens  $A$ 's and  $B$ 's admissible resource regimes so that both will accept common water on the river, replacing  $A$ 's upstream-diversion-only stance with a willingness to accept commons on the shared river. An *exchange-rate table* widens  $B$ 's and  $C$ 's admissible money regimes so that grain-ledger accounting and metal-coin accounting are interconvertible at the trade interface. A *safe-conduct treaty* narrows  $A$ 's and  $C$ 's violence and contract regimes to militia and court, removing the warlord and armed-enforcement options from  $A$  and the monopoly and guild options from  $C$ .

The simulation in `simulation/three_settlement.py` builds the sheaf, enumerates global sections, and computes  $\dim H^0$  and  $\dim H^1$  over  $\mathbb{R}$  by direct rank computation on the coboundary matrix. The full set of repair combinations gives the table below.

Configuration	$ \mathcal{F}(A) $	$ \mathcal{F}(B) $	$ \mathcal{F}(C) $	$\dim H^0$	$\dim H^1$	# global sections
baseline	16	64	32	100	2	0
+ water compact	32	64	32	115	1	0
+ exchange-rate table	16	64	48	116	1	0
+ safe-conduct treaty	4	64	8	70	2	0
water compact + exchange-rate table	32	64	48	132	1	512
water compact + safe-conduct treaty	8	64	8	73	1	0
exchange-rate table + safe-conduct treaty	4	64	12	74	1	0
all three repairs	8	64	12	78	1	512

The baseline admits no coherent macro-order. Two independent edge-disagreement directions remain unresolved: the river (where  $A$ 's upstream water diversion is incompatible with  $B$ 's common and kin-allocated water) and the grain route (where  $B$ 's grain ledger and barter are incompatible with  $C$ 's metal coin and credit). The military road is satisfiable in baseline because  $A$  and  $C$  already share militia violence and court enforcement on their candidate stalks. The road's interface is glued in baseline; its agreement does not depend on any repair. No global section nevertheless exists because all edges have to glue simultaneously, and a single broken edge prevents any consistent assignment regardless of how many other edges are fine.

The three single repairs behave differently from one another. The water compact resolves the river edge, dropping  $\dim H^1$  from 2 to 1. The exchange-rate table resolves the grain edge, with the same

effect. Both leave the section count at zero because the other broken edge still blocks integration. The safe-conduct treaty does something else again. It narrows the violence stalk on  $A$  and  $C$  to militia and the contract stalk to court, which were already satisfiable in baseline through that very intersection. The treaty therefore reduces the size of the cochain spaces (the stalks shrink,  $\dim H^0$  drops from 100 to 70) without reducing  $\dim H^1$ . Repair work directed at an already-glued edge is institutional motion that does not buy compatibility. The Lemma of §2.3 makes the mechanism general: in baseline  $E_b = \{AB, BC\}$ , and the safe-conduct treaty operates on  $AC \in E \setminus E_b$ , so by the contrapositive of the Lemma it cannot enlarge the integer global-section count from zero on its own.

The minimal repair set is the pair {water compact, exchange-rate table}. With both in place, the river and grain edges both glue, the road continues to be satisfiable through its baseline intersection, and 512 admissible global sections emerge. Adding the safe-conduct treaty on top changes neither the section count nor the residual  $\dim H^1$ . In a reconstruction setting where each repair carries a real political cost, this is a load-bearing finding. The safe-conduct treaty looks substantial; it commits two regions to renouncing warlord violence and armed enforcement. It does no compatibility work.

The pattern generalises beyond the toy. Diplomatic effort on edges that are not the obstruction is theatre; diplomatic effort on the obstruction edges is constructive. The model gives the obstruction class as a computable diagnostic, and the diagnostic localises which edges are doing the work. In a larger complex with hundreds of vertices, the same computation ranks candidate repairs by the contribution they make to reducing  $\dim H^1$ . This is the engineering use the framework is designed for.

The 512 admissible sections of the minimum-repair configuration are tuples of local configurations that pass every overlap check, and they are not 512 different macro-orders. Their count factorises transparently. The road forces  $V = \text{militia}$  and  $C = \text{court}$  at both  $A$  and  $C$ , since those are the only intersections of the candidate stalks. The water compact forces  $R = \text{common}$  at both  $A$  and  $B$ . The exchange-rate table forces  $M \in \{\text{grain\_ledger}, \text{metal\_coin}\}$  at both  $B$  and  $C$ , with  $B.M = C.M$ . With those constraints absorbed, the residual freedom on each vertex is small.  $A$  has 2 choices for labor and 2 for knowledge, hence 4 residual configurations.  $B$  has  $2 \cdot 2 \cdot 2 \cdot 2 \cdot 2 = 32$  residual configurations once its money is conditioned on the choice  $C$  makes.  $C$  has 2 resource-rule choices and 2 knowledge choices, hence 4 residual configurations conditioned on its money. The product is  $4 \cdot 32 \cdot 4 = 512$ , and each tuple is a candidate macro-order whose dynamical filter is the sweep of §5.

## 5. Walking the parameter cube

We sample the dynamical layer on an  $11 \times 11 \times 11 \times 11$  grid over  $(v, t, e, l) \in [0, 1]^4$ , 14,641 parameter points in all. At each point we compute the basin distribution over the eight attractors (the Boltzmann distribution over their fitnesses with  $\beta = 12$ ) and assign the dominant attractor to that point. The fraction of the grid won by each attractor is reported as its *basin volume*. The number is a property of the parameter geometry rather than a prediction about historical frequency. It

asks how much of the four-dimensional environment makes each attractor the strongest available choice.

Attractor	Basin volume
Capitalist market	0.031
Guild-merchant	0.101
Kin-commons	0.103
Warlord tribute	0.104
Monastic redistribution	0.109
Commons federation	0.118
State-command	0.184
Neo-feudal	0.249

The capitalist attractor wins 3.1% of the parameter cube. It is the smallest of the eight basins, and it is small because of the geometry of its fitness peak rather than any normative judgement built into the model. Capitalism in this taxonomy peaks where violence is low, transport security is high, energy surplus is high, and literacy retention is high. The corner of the parameter cube where all four conditions meet, after smoothing by the attractor's bandwidth, comes out at three percent. The neo-feudal basin, by contrast, occupies a quarter of the cube, because its fitness peak sits at moderate-to-high violence, low transport, low energy, and moderate literacy, a region the parameter sweep visits much more often than the capitalist corner. State-command takes another eighteen percent because its peak at moderate violence with high transport and high literacy intersects a substantial slab of the cube. Together neo-feudal and state-command win 43.3%.

The corners of the cube reproduce historically recognisable recipes. At  $(v, t, e, l) = (0.1, 0.9, 0.9, 0.9)$ , with low violence, secure transport, abundant energy, and preserved literacy, the capitalist attractor wins 99.7% of the basin mass. This is the early-modern Northwest Europe corner of the parameter cube and it does, as historical intuition would predict, support a capitalist macro-order overwhelmingly. At  $(0.9, 0.1, 0.1, 0.1)$ , with high violence, no transport, no surplus, and no literacy, the warlord-tribute attractor wins 95.7%. This is the post-collapse subsistence countryside in its insecure variant. At  $(0.1, 0.1, 0.1, 0.1)$ , with low violence and low everything else, the kin-commons attractor wins the entire basin mass at 99.99%. This is the post-collapse subsistence countryside in its peaceful variant.

The middle of the cube tells different stories. At  $(0.5, 0.5, 0.5, 0.9)$ , with moderate violence, moderate transport, moderate surplus, and high literacy, the monastic-redistribution attractor wins 96.8%. The corner that capitalist intuition would predict to be inhospitable to a non-market regime is the natural home of monastic and temple redistribution, because the high literacy variable pulls strongly enough toward attractors with literacy-dependent fitness peaks. At  $(0.1, 0.5, 0.5, 0.9)$ , with low violence, moderate transport, moderate surplus, and high literacy, the commons-federation basin wins 99.0%. This is the corner where the Ostrom design principles describe what a polity will tend to assemble. At  $(0.5, 0.9, 0.5, 0.9)$ , with moderate violence,

secure transport, and high literacy, the state-command basin wins 99.2%. This is the corner where bureaucracy beats markets and federations. The model is not silent about anywhere in the cube; it tells a definite story about what kind of macro-order each parameter combination supports.

The borderland between attractors carries its own information. At  $(0.85, 0.15, 0.4, 0.4)$ , with high violence, low transport, moderate surplus, and moderate literacy, the basin splits between neo-feudal (0.683) and warlord tribute (0.315). This is the boundary between the regime that hardens into manorial extraction and the regime that stays in pure protection-rent. The historical transition from late Roman warlord violence to Carolingian manorialism is, in the model's terms, a slow drift along this boundary as the surplus parameter rises and the literacy parameter recovers from below. Scheffer (2009) frames such transitions in social-ecological systems as critical regime shifts driven by slow control variables crossing tipping points; the basin boundaries of the present model give a constructive instance of that framework on institutional rather than ecological dimensions.

A one-dimensional cut through the cube isolates the dynamics relevant to the violence-narrows-the-gluing-space proposition below. Holding  $(t, e, l) = (0.85, 0.85, 0.85)$  at favourable values, with the rest of the cube held capitalist-friendly, and varying  $v$  from 0 to 1 in steps of 0.05, traces the capitalist attractor's basin as violence rises against an otherwise hospitable backdrop.

The basin sits at 0.963 when violence is zero. It stays above 0.9 through  $v = 0.2$ , falls to 0.616 at  $v = 0.4$ , drops to 0.205 at  $v = 0.5$ , 0.084 at  $v = 0.55$ , and 0.042 at  $v = 0.6$ . From there it remains under 0.11 for the rest of the slice. The collapse is sharp and asymmetric, concentrated in the band  $v \in [0.4, 0.6]$ , a  $\Delta v$  of 0.2 across which the basin loses a factor of nearly fifteen.

The basin mass that leaves capitalism does not, in this slice, redistribute over warlord and neo-feudal attractors. With transport, energy, and literacy held high, warlord and neo-feudal attractors remain weak everywhere on the slice; their basin shares stay below 0.11 throughout. The mass migrates almost entirely onto the state-command attractor. State-command's basin rises from 0.058 at  $v = 0.4$  to 0.583 at  $v = 0.5$  and to 0.896 at  $v = 0.65$ , where it dominates the slice. At still higher violence even state-command erodes (its basin falls to 0.279 at  $v = 1.0$ ) and the residual mass spreads thinly across all the other attractors. Violence in the presence of preserved infrastructure drives society toward administered command rather than toward anarchy.

The asymmetry has substantive content. The *collapse* of the capitalist basin and the *destination* of its escaping mass are determined by different parameters. Violence collapses capitalism above a threshold; the literacy and transport parameters then decide which regime captures the released basin mass. Subtract literacy and transport from the slice and the collapsing mass migrates onto warlord and neo-feudal attractors instead. Violence alone does not predict the macro-regime a violence-shocked society falls into.

## 6. Three propositions

The three propositions state what the model entails. None is a new historical claim; each is a consequence of the construction, stated in a form the model can quantify.

**Proposition 1** (Compatibility precedes scale). *Let  $X$  be a connected institutional cell complex and  $\mathcal{F}$  a sheaf of local institutional configurations. A macro-order can scale across  $X$  only if there exists a global section  $s \in \Gamma(X, \mathcal{F})$ , or if institutional repair operations reduce the obstruction class  $o(X, \mathcal{F})$  to zero on the strategically necessary overlaps.*

A civilization that scales is one whose interfaces work. Functional parts in isolation do not assemble into a macro-order on their own. The three-settlement worked example exposes this clearly.  $A$ ,  $B$ , and  $C$  are each internally workable in baseline, and no macro-order exists across them until their water and grain interfaces are reconciled. Bigger versions of the same problem are the daily work of every supranational institution from the Hanseatic League to the European Single Market. The work is identifying which interfaces are blocking integration and engineering the minimal repairs that allow gluing, rather than improving local rules in isolation.

**Proposition 2** (Capitalist reconstruction threshold). *A capitalist global section exists only if the institutional sheaf admits compatible substructures for alienable property, mobile labor, exchange media, contract enforcement, credit assignment, and violence control across a connected trade-producing subcomplex  $X_C \subseteq X$ .*

The empirical content of the proposition is the basin volume reported in §5. Capitalism wins 3.1% of the parameter cube because the simultaneous coincidence of the six conditions listed is geometrically rare, with the rarity following from the conjunction itself rather than from any exotic content in capitalism. The proposition adds nothing to historical knowledge taken individually: economic historians have known for at least half a century that early modern Northwest European capitalism required a particular conjunction of property law, mobile labor, banking infrastructure, secure transport, and political stability. What the model adds is a precise sense of why the conjunction is rare. It is rare because the corner of the parameter cube where all six conditions hold simultaneously is small, and the corner is small because each condition prunes the cube independently and the prunings multiply. Reading the capitalist macro-order as an attractor with a small basin is a different kind of statement from reading it as a contingent historical achievement. The two readings are compatible, with the first one quantifying what the second one only narrates.

**Proposition 3** (Violence narrows the gluing space). *As predation risk increases, the number of feasible global sections decreases, and the system's attractor basins shift toward administered or extractive regimes whose identity depends on the other parameters.*

The violence slice in §5 demonstrates the dynamical version directly. Holding the other three parameters at favourable values, raising violence from  $v = 0.4$  to  $v = 0.6$  collapses the capitalist basin from 0.616 to 0.042. The escaping basin mass falls on the state-command attractor in this slice, because literacy and transport are intact and they pull the mass toward bureaucratic order rather than toward warlord predation. Violence determines that the macro-order will become more administered or more extractive; the standing-stock of literacy and transport determines which. Together with the empirical observation in the original sheaf-theoretic literature (Robinson, 2014) that obstruction classes can be reduced by targeted local interventions, the proposition has practical force. The question facing a reconstruction effort is whether to fight violence before or after

the literacy and transport parameters have collapsed, because the macro-order one ends up with depends on the order of those collapses.

## 7. What the model does not do

The sheaf is necessarily coarse. Seven institutional dimensions with three to five values each cannot capture every distinction that historians draw. The Carolingian villa system, the thirteenth-century Tuscan commune, the late-Tokugawa village federation, and the seventeenth-century Massachusetts Bay town would all classify, in this taxonomy, as variants of two or three of the same attractors. The coarseness is appropriate for the basin-geometry argument the paper makes and inappropriate for the close-grained comparative history those examples deserve. The model claims that the gluing structure on the dimensions it does include is a non-trivial constraint on which macro-orders can exist; it does not claim to predict which specific historical order will arise in any particular place.

The state space is finite by construction. The model identifies which combinations of categorical institutional values glue. It cannot represent continuous gradations within a category, including partial alienability of property, mixed labor regimes within a household, and contract enforcement that runs through a court for some goods and through a guild for others. These gradations are real features of historical institutions; the discrete model trades them for tractability. A continuous extension is possible (each stalk becomes a region of  $\mathbb{R}^7$  with restriction maps as smooth projections) and is left for future work.

Historical calibration is analogue rather than fitting. The eight attractors are stylised. The fitness functions over the parameter cube are transparent bumps chosen to reproduce intuitive historical correspondences (warlord tribute peaks at high violence and low transport; capitalism peaks at low violence and high values of the other three parameters), with no estimation from any dataset of historical regime transitions. The basin volumes reported in §5 are facts about the model's geometry; they carry no predictive weight about the relative historical frequency of macro-regimes. Calibration of the fitness functions against a corpus of post-collapse trajectories (post-Roman West, post-Han China, post-classic Maya, post-Bronze-Age Mediterranean, the various twentieth-century state collapses for which we have detailed records) is a substantial empirical project that would tighten the model's claims; we do not undertake it here.

Obstruction-class computation depends on the chosen restriction maps. The decision that a river edge requires agreement only on the resource dimension, a grain route only on money, and a road on violence and contract, is a modelling decision with real content. A modeller who treats a river as carrying agreement on all seven dimensions will get a different obstruction class on the same cell complex. The point of the framework is to make those decisions explicit and to let the reasoning about them be a public part of the model rather than a hidden assumption. The model is responsive to its restriction maps; it is fragile to wholesale replacement of them.

The capitalist attractor carries no normative weight in the analysis. The model treats it as one attractor among eight. The finding that its basin volume is 3.1% is descriptive, following from

the geometry of the four-parameter cube together with the fitness specifications. Whether a 3.1% basin is “small” or “large” depends on the reader’s prior; the model itself is silent on the question. The framework can equally be used to ask under which interventions the commons-federation basin grows, or the monastic basin, or the state-command basin. None of these questions has a privileged status in the formalism.

A final and harder limit is that the model’s choice of *what counts as a region* carries content. Settlements are vertices because the model needs vertices to be the loci of institutional decision; historical settlements are not always sharp. Diasporic networks, transhumant pastoralists, and modern remote workforces complicate the simple vertex-as-place picture. A more elaborate model would allow vertices to overlap or to be inhabited by multiple institutional populations with their own internal interfaces. The version of the model presented here treats vertices as discrete; its conclusions transfer to richer models only if the gluing structure they impose between *populations* in the same place is structurally similar to the gluing structure imposed between *places* in this paper.

## 8. Engineering interfaces

Civilizations are rebuilt from local rule systems outward, with the rules at the seams between regions deciding whether a coherent macro-order can stand. Sheaf cohomology gives the formal language for the compatibility condition. Attractor dynamics gives the formal language for the stability filter that selects among the orders the compatibility allows. The two layers together describe what a reconstruction effort actually faces.

The model’s most concrete deliverable is a diagnostic. Given a cell complex of settlements and a description of their institutional repertoires, the model identifies which interfaces are blocking integration and which candidate repairs would actually move the obstruction class. The three-settlement example demonstrates the diagnostic in a setting small enough to inspect: of three repairs that look comparably substantial in political terms, only two do compatibility work. The third is institutional theatre. In a larger landscape the same computation ranks candidate repairs by the contribution they make to compatibility; a reconstruction programme that proceeds in obstruction-reducing order differs sharply from one that proceeds in cosmetic order, and the model exposes the difference.

The basin geometry is the model’s other principal deliverable, and the harder one to live with. The capitalist attractor is the smallest of the eight basins because it requires the simultaneous coincidence of four parameter conditions that are individually common and jointly improbable. The neo-feudal and state-command basins together occupy more than two-fifths of the parameter cube. The basin geometry carries no judgement about whether this is good or bad. It carries the structural fact that the macro-order most contemporary politics treats as the natural fallback of any rebuilding society is, in the model’s terms, an unusual coincidence rather than a default. A reconstruction effort that aims at the capitalist attractor without attending to the parameter conditions that make its basin reachable is aiming at a target whose basin its actions are not enlarging.

The three propositions distil the substantive content. Compatibility precedes scale; the capitalist

macro-order requires the simultaneous gluing of six dimensions across a trade-producing subcomplex; violence narrows the gluing space and the destination of the basin mass that leaves capitalism is determined by what other parameters survive. None of the three propositions is novel as a statement; their contribution is that they emerge from a model in which they can be quantified, and where the parameter conditions under which each holds or fails can be examined directly. The model's coarseness is a deliberate feature of this generation of the work; the next generation will need finer dimensions, continuous stalks, and historical calibration of the attractor fitness functions. None of those extensions changes the structural argument that civilizations are local-to-global problems before they are anything else, and that their reconstruction is, accordingly, the engineering of interfaces under environmental conditions that the engineers do not fully control.

## References

- Abramsky, S., & Brandenburger, A. (2011). The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13(11), 113036.
- Acemoglu, D., Johnson, S., & Robinson, J. A. (2001). The colonial origins of comparative development: An empirical investigation. *American Economic Review*, 91(5), 1369–1401.
- Acemoglu, D., & Robinson, J. A. (2012). *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. Crown.
- Arthur, W. B. (1994). *Increasing Returns and Path Dependence in the Economy*. University of Michigan Press.
- Bodnar, C., Di Giovanni, F., Chamberlain, B. P., Liò, P., & Bronstein, M. M. (2022). Neural sheaf diffusion: A topological perspective on heterophily and oversmoothing in GNNs. In *Advances in Neural Information Processing Systems* 35.
- Cline, E. H. (2014). *1177 B.C.: The Year Civilization Collapsed*. Princeton University Press.
- Curry, J. (2014). *Sheaves, Cosheaves and Applications*. PhD thesis, University of Pennsylvania.
- Diamond, J. (2005). *Collapse: How Societies Choose to Fail or Succeed*. Viking.
- Edelsbrunner, H., & Harer, J. L. (2010). *Computational Topology: An Introduction*. American Mathematical Society.
- Folke, C. (2006). Resilience: The emergence of a perspective for social-ecological systems analyses. *Global Environmental Change*, 16(3), 253–267.
- Ghani, N., Hedges, J., Winschel, V., & Zahn, P. (2018). Compositional game theory. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, 472–481.
- Greif, A. (2006). *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge University Press.
- Hansen, J., & Gebhart, T. (2020). Sheaf neural networks. *NeurIPS Workshop on Topological Data Analysis and Beyond*.

- Hansen, J., & Ghrist, R. (2019). Toward a spectral theory of cellular sheaves. *Journal of Applied and Computational Topology*, 3(4), 315–358.
- Heather, P. (2006). *The Fall of the Roman Empire: A New History of Rome and the Barbarians*. Oxford University Press.
- Holling, C. S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4, 1–23.
- Lewis, M. E. (2009). *China Between Empires: The Northern and Southern Dynasties*. Harvard University Press.
- McAnany, P. A., & Yoffee, N. (Eds.). (2010). *Questioning Collapse: Human Resilience, Ecological Vulnerability, and the Aftermath of Empire*. Cambridge University Press.
- North, D. C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge University Press.
- North, D. C., Wallis, J. J., & Weingast, B. R. (2009). *Violence and Social Orders: A Conceptual Framework for Interpreting Recorded Human History*. Cambridge University Press.
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.
- Riess, H., & Ghrist, R. (2022). Diffusion of information on networked lattices by gossip. In *Proceedings of the 61st IEEE Conference on Decision and Control*, 5946–5952.
- Robinson, M. (2014). *Topological Signal Processing*. Springer.
- Scheffer, M. (2009). *Critical Transitions in Nature and Society*. Princeton University Press.
- Strogatz, S. H. (2015). *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering* (2nd ed.). Westview Press.
- Tainter, J. A. (1988). *The Collapse of Complex Societies*. Cambridge University Press.
- Tilly, C. (1990). *Coercion, Capital, and European States, AD 990–1990*. Blackwell.
- Turchin, P. (2003). *Historical Dynamics: Why States Rise and Fall*. Princeton University Press.
- Walker, B., Holling, C. S., Carpenter, S. R., & Kinzig, A. (2004). Resilience, adaptability and transformability in social-ecological systems. *Ecology and Society*, 9(2), 5.
- Wickham, C. (2005). *Framing the Early Middle Ages: Europe and the Mediterranean, 400–800*. Oxford University Press.
- Williamson, O. E. (2000). The new institutional economics: Taking stock, looking ahead. *Journal of Economic Literature*, 38(3), 595–613.