

Humanized Autonomous Organization

1. Introduction
2. Structure
 1. Structural Overview of the HAO Framework
 2. United Micro Enterprises (UMEs)
 3. Strategic Enterprise Partnerships (SEPs)
 4. Micro Enterprise Ecosystem (MEE)
 5. Public Market Interfaces (PMIs)
3. Governance
 1. Membership and Quorum Rules
 2. Voting and Expertise Weighting
 3. Accountability Mechanisms
 4. Cross-Unit Governance Instruments
4. Economics
 1. Trickle-Out Investment Flow
 2. Revenue Allocation Framework
 3. Strategic Enterprise Partnership (SEP) Economic Agreements
 4. Equity and Member Compensation Models
 5. Reinvestment Mechanisms
 6. External Interface Economics and Investor Buffers
5. Lifecycle
 1. UME Lifecycle and Transition States
 2. Strategic Enterprise Partnership (SEP) Lifecycle and Governance
 3. Enterprise Collapse and Containment Protocols
 4. Governance Protocol Evolution
 5. Intentional Ephemerality
 6. Lifecycle Variants by Unit Type
6. Human Systems
 1. Socioemotional Safety and Shared Language
 2. Cultural Onboarding and Ritual Design
 3. Conflict Engagement as Organizational Literacy
 4. Regenerative Learning Systems
 5. Synthesis — Human Systems as Infrastructure
7. Legal
 1. Modular Legal Forms and Multi-Level Structures
 2. Jurisdictional Flexibility and Nesting
 3. Commons-Based Intellectual Property and Licensing Models
 4. Public Market Interaction Framework
8. Adversarial Analysis
 1. Threat Taxonomy
 2. Governance Capture
 3. Framework Capture
 4. External and Technical Vectors

5. [Polycentric Structural Failure](#)
6. [Countermeasures and Detection](#)
7. [Open Problems](#)
9. [Evaluation](#)
 1. [Metrics Across Capitals](#)
 2. [Participation Quality and Alignment Audits](#)
 3. [Continuous Sensing and System Adaptation](#)
 4. [AI-Augmented Governance Monitoring](#)
 5. [Systemic Learning and Evolution Protocols](#)
10. [Deployment](#)
 1. [Seeding an HAO Network](#)
 2. [provide.io as a Deployment Engine for HAO](#)
 3. [Genesis Process for the First UMEs](#)
 4. [Strategic Enterprise Partnership Structuring](#)
 5. [Iterative Scaling Phases](#)
 6. [Patterns of Interoperability](#)
 7. [Formation Evidence](#)
11. [ICN](#)
 1. [Structural Features: UMEs, SEPs, MEE, and HAO](#)
 2. [Governance Framework and DEA Protocols](#)
 3. [Economic Model: Trickle-Out Capital Allocation, Diminishing Contributions, and Equity](#)
 4. [Technical Stack and Platform Design](#)
 5. [Public Market Interfaces \(PMIs\)](#)
12. [Member Trust Union](#)
 1. [Multi-Level Trust Architecture](#)
 2. [Financial Product Design in the Member Trust Union \(MTU\)](#)
 3. [Governance and Risk Distribution in the Member Trust Union \(MTU\)](#)
 4. [Technology and Physical-Digital Integration in the MTU](#)
 5. [Member Experience and Compensation Systems in the MTU](#)
 6. [Integration with ICN and External Markets](#)
 7. [Ethical and Regulatory Anchoring](#)
 8. [MTU as a Domain-Specific HAO Implementation](#)
 9. [Physical Branch Model](#)
 10. [Jurisdictional Compliance](#)
 11. [Legal Vehicle Options](#)
13. [Advanced](#)
 1. [Cross-HAO Interoperability and Network Federation](#)
 2. [AI-Augmented Governance and Decision Support](#)
 3. [Trust Algorithms and Verification Frameworks](#)
 4. [Mutual Credit Systems and Internal Economies](#)
 5. [Bioregional and Sector-Specific HAO Implementations](#)
 6. [Systems Simulation, Modeling, and Stress Testing](#)
14. [Appendices](#)
 1. [Appendix: Visual Concepts for Section 13.5 – Member Experience and Compensation Systems](#)
 2. [!\[\]\(694fcb4611893e9db5249daba48abfc1_img.jpg\) Worker-Owned Café \(Portland, OR\): UME Viability Models](#)
 3. [!\[\]\(8ec8d5dc48934930a762fecf6ecbe179_img.jpg\) Updated Models: Customer Volume & Order Value Integration](#)
 4. [Citation Audit](#)
 5. [HAO Lineage and Vocabulary Genealogy](#)
 6. [HAO Open Questions Register](#)

7. HAO Corpus References
8. Research Evidence Map
9. HAO Source Integration Register
10. HAO Framework Terminology (Extended Glossary)

1.0 • Introduction

1.1 Purpose and Research Context

This document describes a systems-level architecture for a new organizational model: the Humanized Autonomous Organization (HAO) — the network’s coordinating framework. The HAO is intended to support member well-being in conditions of increasing complexity, automation, and global interdependence. It is a full-stack socio-technical framework, grounded in theory and specified for implementation.

The goal is to define the model clearly enough to be replicated, adapted, and evolved across domains: economic networks, financial cooperatives, local service ventures, globally distributed ecosystems.

Two reference implementations are explored in depth:

- The **Integrated Cooperative Network (ICN)** — the reference cooperative business network: a network of small, semi-autonomous business units called UMEs (small, self-managing venture teams) that generate value in local and global markets.
- The **Member Trust Union (MTU)** — the network’s credit-union-like financial institution: a trust-centered, polycentric financial system structured as an alternative to traditional banking and fintech.

This paper articulates the shared architecture that underlies these systems and offers the tools, language, and design principles needed to build others.

1.2 Limitations of Traditional Organizational Models

Since the industrial era, most organizations share a common structure: hierarchical control, centralized ownership, and efficiency-driven workflows, producing large gains in scale along with the costs listed below.

Challenges include:

- Disconnection between those who create value and those who capture it
- Organizational brittleness in the face of complexity and disruption
- Misalignment between financial incentives and human or ecological well-being
- Cultural stagnation inside rigid systems not designed for learning or inclusion

Even newer models such as DAOs, platform co-ops, or remote-first startups often carry forward these assumptions with newer tools.

HAOs differ in structure: human needs are the starting point, autonomy is distributed rather than centralized, and technology supports resilience rather than control.

1.3 Overview of the HAO Framework

The Humanized Autonomous Organization is organized around the following design premises:

- **Human primacy:** Technology supports human judgment rather than replacing it.
- **Distributed autonomy:** Power and decision-making are shared across the network rather than concentrated in a single unit.
- **Value alignment:** The organization’s structure, actions, and incentives reflect shared principles.
- **Resilience over efficiency:** Redundancy, adaptability, and trust networks are prioritized.

Structurally, the HAO is composed of:

- A core coordinating system (the HAO layer)
- A network of autonomous operational units (UMEs)
- Governance frameworks that evolve over time: the DEA (a versioned operating agreement replacing fixed bylaws) and the AGF (the network’s layered governance system)
- Economic protocols that reward value creators and protect long-term integrity
- Systems for culture, trust, and conflict resolution at scale

Where DAOs rely on code to enforce structure, HAOs treat governance, learning, and social trust as first-class design elements alongside code and protocol.

1.4 HAO Compared to Existing Models

To situate the HAO, it helps to compare it with common organizational patterns:

Model	Key Characteristics	Limitations
Traditional Corporations	Centralized control, linear workflows, shareholder primacy	Single point of failure; profit flows to shareholders before operating units; slow to adapt

Model	Key Characteristics	Limitations
Platform Cooperatives	Member-owned platforms, democratic control	Struggle with scale, often lack technical integration
DAOs	Smart contract-based governance, global coordination	Token-weighted voting excludes qualitative judgment; poor conflict handling; unclear value alignment
HAOs	Human-centric, polycentric, trust-based with tech augmentation	Higher design complexity; requires cultural onboarding
The HAO framework combines ethical infrastructure, economic redesign, and adaptive governance.		

1.5 Reader Orientation and Scope

This paper is written for system designers, founders, technologists, organizers, policy architects, and researchers evaluating alternatives to centralized or purely algorithmic organizational models.

It offers:

- A breakdown of HAO’s architecture, from structure to compensation
- Implementation blueprints for the ICN (economic) and MTU (financial)
- A shared terminology for theory and practice
- Diagrams, governance templates, and lifecycle protocols
- Guidance for implementation, evaluation, and scaling

The aim is clarity: the framework specifies the model in enough detail to check an implementation against it.

2.0 • Structure

§2.1 Structural Overview of the HAO Framework

The Humanized Autonomous Organization (HAO) — the network’s coordinating framework — is a pattern language for organizing human and technological systems around shared values, distributed autonomy, and mechanisms for reinvesting economic surplus across the network.

This section describes the primary structural components of the HAO, which together form a modular, composable system. These structures are not rigid: they provide the scaffolding for culture, trust, governance, and value flow.

The five foundational elements explored here are:

- **The HAO layer:** A coordination and systems integration layer that maintains integrity across a distributed network without centralizing power.
- **United Micro Enterprises (UMEs)** — small, self-managing venture teams of up to ~15 people: the autonomous, semi-permeable value-generating nodes of the system.
- **Strategic Enterprise Partnerships (SEPs)** — joint ventures between teams: temporary or long-term collaborations between UMEs to pursue aligned missions.
- **Micro Enterprise Ecosystem (MEE)** — the network’s protected internal economy: a protected yet porous economic and cultural environment that allows UMEs to operate while maintaining alignment with core principles.
- **Public Market Interfaces (PMIs)** — buffer companies between the network and outside investors: specialized structures that allow selective engagement with external capital or market forces without compromising internal governance.

Each component is meaningful on its own and interdependent with the others, forming a system intended to grow, self-correct, and evolve while maintaining coherence.

We begin with the HAO itself.

2.1 The Humanized Autonomous Organization (HAO)

Every HAO-based system includes the HAO layer: a coordinating, enabling, and value-protecting structure that facilitates communication, coherence, and shared learning across nodes, rather than functioning as a command center, corporate headquarters, or founder-led hub.

Purpose of the HAO

The HAO exists to:

- Maintain alignment between distributed units and shared principles
- Provide enabling infrastructure: governance, technical, legal, financial
- Ensure interoperability between local autonomy and global coherence
- Facilitate cross-network coordination, learning, and resilience
- Protect the system’s integrity against value extraction outside agreed terms or misaligned incentives

Its role is to make agency and coherence compatible at scale without becoming bureaucratic or controlling.

Characteristics

The HAO is defined by a few key structural features:

- **Non-centralized but coordinating:** It does not “own” the network, but provides shared infrastructure and arbitration mechanisms.
- **Constitutional Layer:** It maintains the Dynamic Enterprise Agreement (DEA) — a versioned operating agreement replacing fixed bylaws — that underpins the system’s governance.
- **Polycentric Facilitation:** It enables multiple centers of decision-making across contexts and domains, using protocols like the Adaptive Governance Framework (AGF) — the network’s layered governance system.
- **Implementation-Agnostic:** The HAO model can be implemented by different entities (e.g., provide.io) depending on context, domain, and regional needs.
- **Temporal Role:** In some networks, the HAO may phase out over time, with responsibility distributed fully to autonomous entities.

HAO Functions

Core HAO responsibilities include:

- **Governance Enablement:** Managing the lifecycle and integrity of the DEA and AGF across the network.
- **Tech Infrastructure Deployment:** Provisioning and maintaining distributed ledger systems, reputation layers, and cultural monitoring tools.
- **Capital Coordination:** Managing trickle-out investment flows, revenue reconciliation, and reinvestment protocols.
- **Conflict Mediation:** Offering neutral third-party resolution mechanisms and shared socioemotional language systems.
- **Standards and Integration:** Ensuring compatibility across protocols, data formats, and organizational lifecycles.

HAO is Not:

- A CEO or “founder” structure
- A traditional holding company or central bank
- A blockchain DAO governed solely by token-based consensus
- A static or permanent authority
- An outside funder extracting profit from edge contributors

It functions as a systems integrator that evolves over time, structured as a coordinating pattern rather than a corporate entity.

Multiple Instantiations

There is no single “correct” HAO. In one network, it might be implemented by a systems integration entity like *provide.io*. In another, it could emerge as a shared cooperative council. What matters is adherence to structural roles and principles, not branding.

As long as it fulfills the responsibilities of coordination, coherence, and integrity, without concentrating power or value, the HAO pattern is considered valid.

§2.2 United Micro Enterprises (UMEs)

United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people — are the fundamental building blocks of a system coordinated by the HAO (the network’s coordinating framework). Each UME is a semi-autonomous, human-scaled unit responsible for generating value, stewarding resources, and participating in the broader network. UMEs operate independently while remaining bound by shared agreements, technologies, and ethical principles.

Purpose and Role

The UME exists to:

- Create localized or domain-specific value: a product, service, knowledge artifact, or cultural practice
- Sustain livelihoods for members through compensation and participatory decision-making
- Experiment, adapt, and evolve within the boundaries of shared principles
- Interface with peers and with the HAO through defined protocols rather than a subordinate relationship
- Provide members a shared cultural and organizational context for their work

Most economic activity, social learning, and organizational identity in the network occur within UMEs rather than in the HAO layer.

Design Characteristics

UMEs share a structural and philosophical architecture, but remain individually expressive:

- **Human-scale:** Designed for internal cohesion, shared context, and participatory processes. Most UMEs cap active members to maintain sociological coherence.
- **Semi-permeable:** Open to collaboration and exchange, but with the ability to set boundaries around internal decisions and identity.
- **Self-governing:** Governed by a localized instantiation of the DEA (a versioned operating agreement replacing fixed bylaws) and AGF (the network’s layered governance system), adapted to the UME’s specific domain, language, and pace.
- **Multi-capital aware:** Oriented around financial, cultural, social, ecological, and intellectual value.
- **Lifecycle-bound:** Born through intentional formation processes and retired, split, or recombined through structured degradation protocols.

Operational Structure

Most UMEs function around the following components:

- **Circle or Crew:** The core operating group responsible for the day-to-day mission.
- **Lead Steward(s):** Rotational facilitation roles focused on coordination rather than command or management.
- **Governance Layer:** A local version of the AGE, supporting role assignment, feedback, and consent-based decision-making.
- **Accounting/Value Ledger:** Tracks time, capital, commitments, and performance against both financial and non-financial KPIs.
- **Participation Protocol:** Describes how new members enter, level up, or exit; includes contribution thresholds, rights, and trust-building stages.

While structures may vary, they remain legible to the network: a UME can interface with others without translation layers or opaque processes.

Autonomy and Interdependence

UME autonomy is bounded by the network’s shared agreements, allowing independence without fragmentation.

- UMEs set their own mission, but align with broader ethical principles (e.g. the ETHICAL framework).
- They manage their own cash flow, but contribute agreed percentages back to the HAO.
- They engage in contracts and SEPs (joint ventures between teams), but within the terms of the Dynamic Enterprise Agreement.
- They define their own internal culture, but are accountable to the shared culture monitoring and trust verification mechanisms.

This structure allows UMEs to experiment, express local variation, and act with agility, while also participating in a larger, resilient system.

Emergence and Dissolution

Each UME is created through a structured genesis protocol, often initiated by a group of aligned individuals or seeded via HAO investment. New UMEs must pass a readiness assessment, which includes:

- Strategic alignment
- Team cohesion
- Resource plan
- Trust and culture foundations
- Governance readiness

UMEs can be:

- **Retired** (graceful dissolution due to lifecycle end)
- **Absorbed** (if their work is continued by a larger UME)
- **Split** (into two new UMEs after growth or divergence)
- **Rebooted** (if temporarily paused or restructured)

Degradation and exit are treated as stages in the UME lifecycle rather than as failures.

UME ≠ “Small Business”

Though they may resemble small enterprises in scale, UMEs are not defined by market logic alone. They are:

- Embedded in the network’s ethical and governance systems
- Tied into shared value tracking and redistribution
- Cultivating cultural, social, and ecological capital

UMEs function as interdependent parts of a larger system rather than as siloed ventures pursuing isolated outcomes.

§2.3 Strategic Enterprise Partnerships (SEPs)

United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people — act as autonomous units of value creation. Strategic Enterprise Partnerships (SEPs) — joint ventures between teams — are how UMEs collaborate, formally, temporarily, or semi-permanently, to pursue goals beyond an individual UME’s scope. A SEP is not a merger, acquisition, or hierarchical relationship; it is a collaboration framework defined by agreements, shared ownership, and bounded scope.

SEPs allow the network coordinated by the HAO (the network’s coordinating framework) to scale horizontally without centralization and without sacrificing local autonomy.

Purpose and Function

SEPs exist to:

- Enable coordinated action across multiple UMEs without dissolving their autonomy
- Pursue shared missions or access new markets that are impractical for individual UMEs to reach alone
- Pool resources, capabilities, or infrastructure in ways that benefit all parties
- Prototype or pilot initiatives that may eventually become new UMEs or system-wide features
- Provide bounded accountability, legal clarity, and economic coordination for joint ventures

In short, SEPs let UMEs work together without forming a new bureaucracy or permanently binding themselves to one another.

Key Characteristics

A Strategic Enterprise Partnership is defined by several core design traits:

- **Collaborative, not Competitive:** Formed around mutual benefit and aligned mission.
- **Scoped and Time-bound:** Defined by a charter, with milestones and review points.
- **Legally Structured:** Can be instantiated as a joint LLC, contract-based collaboration, or trust-bound cooperative venture.
- **Governance-Aware:** Uses a hybrid of the partners' local governance systems, with agreed fallback protocols for conflict or impasse.
- **Dynamic Equity Participation:** Revenue or asset rights are distributed based on contributions and commitments, not fixed shares.

Each SEP is unique, but all share a common structural grammar that keeps them legible to the network.

Formation Process

The SEP lifecycle is governed by a **Shared Project Formation Protocol**, which includes:

1. **Intent Declaration:** Public or private articulation of the mission, timeline, partners, and expected outcomes.
2. **Partner Vetting and Alignment:** Capability assessment, value alignment review, and legal-technical readiness.
3. **SEP Charter Drafting:** Agreement on scope, decision protocols, value flows, and conflict resolution.
4. **Resource Commitment:** Each UME (and potentially the HAO) declares what it is contributing: time, capital, access, tech, reputation.
5. **Legal Instantiation** (if needed): Setup of joint legal entity or binding contract structure.
6. **Launch:** Initiation of operations, tracked against shared KPIs and milestones.

SEPs can be created quickly (some form and deploy within days using standardized templates), while others evolve over months for more complex initiatives.

Governance Structure

SEP governance is designed to be:

- **Lightweight:** Focused on coordination rather than control.
- **Proportional:** Voting power or influence can be based on stake, effort, or domain expertise.
- **Fail-safe integrated:** Includes predefined fallback paths if trust breaks or partners exit.
- **Versioned:** SEP charters are treated as living documents and updated through formal processes.

A SEP uses facilitated coordination, with rotating roles such as:

- **SEP Steward:** Maintains alignment and rhythm
- **Operations Lead:** Handles execution-level oversight
- **Finance Contact:** Manages pooled budgets and revenue sharing
- **Governance Contact:** Responsible for protocol adherence and internal checks

Roles are filled by partner UMEs, often rotated or reassigned periodically.

Economic Agreements

SEP value flows are encoded in a **Dynamic Revenue Participation Model**, which defines:

- **Revenue Distribution Rules:** How proceeds are split across participating UMEs (and potentially the HAO)
- **Expense Responsibility:** Who bears what operational or capital costs
- **Equity and IP Rights:** Whether the output is shared, licensed, or owned by a specific UME
- **Reinvestment Pooling:** Option to allocate a % of returns to further SEP expansion or spinouts

SEP agreements are structured to change as partners' contributions or circumstances shift.

Examples

Example 1: Multi-UME Product Launch

Three UMEs—one specializing in design, one in software development, and one in education—form a SEP to create an ethical edtech platform. They sign a SEP charter, launch a joint pilot, and agree to distribute revenue proportionally based on effort and resource inputs.

Example 2: Regional Market Interface

A cluster of UMEs in a geographic region form a SEP to interface with a local logistics provider. They create a pooled legal entity, negotiate as a collective, and rotate service contracts across members.

Example 3: Public Infrastructure Build

The HAO initiates a SEP with five mature UMEs to co-develop a new governance tool. Contributions include funding, development time, and operational testing environments. All resulting infrastructure is licensed back to the network.

Exit and Transition

SEPs are designed to end, evolve, or spin out. At any point:

- Partners can exit under defined protocols
- The SEP can end upon completion of its mission
- A SEP can spin out into a new UME, SEP, or shared infrastructure layer
- If conflict emerges, fallthrough governance and mediation procedures are followed, with HAO mediation only as a last resort

SEP dissolution does not indicate failure; it indicates that the SEP's bounded intent has reached a natural conclusion or transformed into something else.

Relationship to UMEs and HAO

SEPs sit at the meso-layer between UME autonomy and HAO coherence:

- They are formed by UMEs
- Often use HAO infrastructure
- Can be funded, seeded, or incubated by the HAO
- And may eventually inform broader network evolution

They are a mechanism for scaling the network without centralizing it.

§2.4 Micro Enterprise Ecosystem (MEE)

The Micro Enterprise Ecosystem (MEE) — the network's protected internal economy — is the environment within which UMEs (small, self-managing venture teams) and SEPs (joint ventures between teams) operate. It regulates the flow of information, capital, and technology among the autonomous units of the network coordinated by the HAO (the network's coordinating framework), while managing controlled interfaces with external markets. The MEE provides a buffered space intended to support innovation and local adaptation without compromising network-wide cohesion.

Purpose and Function

The primary functions of the MEE include:

- **Shielding UMEs:** Creating a protected domain where emerging and mature UMEs can operate with reduced exposure to disruptive external market forces.
- **Facilitating Collaboration:** Providing a common ground for diverse UMEs to exchange resources, share best practices, and form SEPs.
- **Maintaining Cultural Coherence:** Reinforcing shared values and ethical standards across the network, so that every UME remains aligned with the overarching HAO principles.

- **Controlled Market Engagement:** Allowing selective, measured interaction with external economic systems via Public Market Interfaces (PMIs) — buffer companies between the network and outside investors — while preserving internal autonomy.
- **Resource Allocation:** Acting as the conduit for cross-enterprise investments, reinvestment mechanisms, and mutual aid protocols.

In essence, the MEE is the ecosystemic “soil” that sustains the network, balancing protection with permeability.

Structural Characteristics

The MEE is built upon several key characteristics:

- **Protected Boundaries:** The MEE is demarcated by defined protocols that regulate the flow of information, capital, and technology. These “boundaries” are intentionally semi-permeable; they allow necessary external inputs while maintaining a core of shared values and practices.
 - **Cultural and Operational Homogeneity:** Although UMEs retain individual identities, the MEE establishes a baseline for operational language, ethical norms, and performance metrics. This framework supports common operational expectations across participants.
 - **Layered Interfacing:** Internally, the MEE is organized into sub-domains: clusters of UMEs that share similar missions or market domains. Externally, it connects to Public Market Interfaces (PMIs) that serve as controlled gateways to broader economic activity.
 - **Feedback and Adaptation Loops:** The ecosystem is dynamic, with built-in mechanisms for continuous learning and feedback. Regular audits, cultural assessments, and resilience checks help the MEE adapt to internal evolution and external disruptions.
 - **Distributed Resource Pools:** Resources (financial, intellectual, or social) are distributed throughout the MEE rather than held by a central entity, supporting interdependence and rapid reallocation in response to opportunities or challenges.
-

Governance and Management within the MEE

Although the MEE is not a centralized authority, it is governed by a series of network-level protocols that maintain system integrity:

- **Common Protocols:** Shared governance documents, including the Dynamic Enterprise Agreement (DEA) — a versioned operating agreement replacing fixed bylaws — and the Adaptive Governance Framework (AGF) — the network’s layered governance system — extend into the MEE, so that each UME and SEP adheres to common standards.
- **Cultural Monitoring Systems:** Tools such as Value Alignment Monitoring (VAM) — ongoing checks that actions match stated principles — and Enterprise Culture Cultivation (ECC) are used across the ecosystem. They measure adherence to core values and flag divergences before they become systemic.
- **Boundary Management:** Specialized sub-teams or committees (often emerging organically from within the network) monitor external interfaces. They ensure that external pressures, such as market volatility or regulatory changes, are absorbed and managed without fracturing internal cohesion.
- **Resource Reinvestment Protocols:** The MEE facilitates structured reinvestment of network resources. When UMEs generate surplus, a portion is pooled and redistributed across the ecosystem to nurture under-resourced units or fund joint projects.

MEE governance combines bottom-up initiatives with top-down coordination through the HAO layer.

Operational Dynamics

The everyday operations within the MEE follow these core dynamics:

1. **Information Flow and Transparency**
 - Open channels of communication allow UMEs and SEPs to share insights, performance data, and cultural narratives.
 - Regular feedback sessions and cross-enterprise reviews keep the ecosystem aligned with its ethical and operational standards.
 2. **Adaptive Resource Sharing**
 - Resources (capital, technology, mentorship) are dynamically allocated based on real-time needs and opportunities.
 - Peer-to-peer micro-loans, shared asset financing, and cooperative procurement strategies are commonplace.
 3. **Boundary Reinforcement and Flexibility**
 - The MEE’s semi-permeable boundary is continuously managed to optimize the balance between insulation and external connectivity.
 - This involves protocols for “soft entry” of external partners and “hard exit” mechanisms when UMEs deviate significantly from network norms.
 4. **Resilience Building**
 - Regular drills and scenario planning test the ecosystem’s robustness against shocks.
 - The MEE is designed to be antifragile: it strengthens from disruptions rather than being weakened by them.
-

Case Illustration

Consider a regional cluster of UMEs operating in a cooperative urban agriculture network. Within the MEE:

- **Protected Spaces:** These UMEs share land, water, and distribution channels, insulated from volatile commodity markets.
- **Collaborative Protocols:** They use standardized practices for crop rotation, resource pooling, and agreed-upon revenue sharing, documented in a common operational handbook.
- **Adaptive Interfaces:** A local Public Market Interface is established for direct-to-consumer sales, intended to maintain the cooperative’s stated values (sustainability, community focus) as it engages with the broader market.
- **Resilience Checks:** Periodic community audits and environmental impact assessments check whether the ecosystem meets its stated environmental and operational targets.

This illustration shows the MEE functioning as both a protective envelope and an active facilitator of growth, collaboration, and innovation.

Conclusion

The Micro Enterprise Ecosystem is the environment in which the HAO’s autonomous units operate. By establishing clear, adaptive boundaries, reinforcing shared cultural and operational standards, and enabling dynamic resource flows, the MEE is intended to support a network that is resilient and responsive to change.

The next section covers Public Market Interfaces (PMIs): the structured interfaces through which the network coordinated by the HAO engages with external economic forces while preserving internal governance.

§2.5 Public Market Interfaces (PMIs)

The **Public Market Interface (PMI)** — a buffer company between the network and outside investors — is the mechanism by which a Humanized Autonomous Organization (HAO) — the network’s coordinating framework — engages with external markets, investors, and regulators while preserving internal governance.

PMIs are semi-autonomous economic intermediaries that convert between internal accounting logic and external ROI-based logic. They are neither full network members nor conventional for-profit shells, but purpose-specific interfaces that manage friction and mismatch between the two.

Purpose and Function

The PMI exists to:

- Facilitate value exchange between the ICN (the reference cooperative business network) and the outside world (e.g., customers, capital providers, governments, platforms).
- Limit UMEs’ (small, self-managing venture teams) direct exposure to volatile or misaligned external economic incentives.
- Offer investor on-ramps that do not require full conversion to standard equity models.
- Create legal, financial, and regulatory wrappers for market-facing activities.
- Translate between internal operating principles and commercial norms.

This lets the PMI handle compliance, branding, and ROI interface logic without altering internal governance.

Core Characteristics

A Public Market Interface typically includes:

Trait	Description
Externally facing	Designed to interact with consumers, clients, investors, and regulators.
Mission-buffered	Holds the structural boundaries on behalf of the network.
Flexible in form	Can take the shape of a public-benefit LLC, coop-corporation hybrid, etc.
HAO-governed	Bound by agreements with the HAO to ensure value alignment and reinvestment.
Revenue-transformative	Converts profits or ROI expectations into network-compatible flows.
Value-aligned branding	Represents the network’s stated values in public communications.

Implementation Options

PMIs are modular and contextual. Their structure depends on domain, market, and legal environment.

A. Contribulo-Type Entity (Example PMI)

As a reference, the example of **Contribulo** (structured in the model as a hybrid LLC) illustrates a potential PMI:

- Structured as a hybrid LLC
- Ownership allocation and any outside-investor share remain pending an owner decision
- Revenue streams include licensing, consulting, or product resale
- In the model, Contribulo remits part of profits to the HAO for reinvestment in the ecosystem
- Holds brand rights, public-facing web properties, and compliance liabilities

This provides a bounded channel for external capital into internal value production.

B. CoopCycle-Like Licensing Wrapper

Another PMI option is a commons-based license steward:

- Holds intellectual property under free/libre licensing conditions
- Grants usage rights only to UME-aligned entities or federated partners
- Manages enforcement and interface with courts, governments, or corporations
- Enables rapid deployment of infrastructure (e.g., logistics, data networks)

This model offers a path to scale without venture capital.

C. Regional Market Shell

A PMI can also take the form of a regional shell entity that:

- Aggregates multiple UMEs in a territory
- Handles compliance, import/export, VAT, tax reporting
- Offers unified contract and insurance layers
- Runs front-office operations (e.g., customer support, unified billing)

This lets participants operate within regulated systems without navigating that complexity individually.

Economic Role

PMIs support a dual-currency dynamic:

1. External Economic Translation

- External revenue, equity, or investment is received in conventional terms.
- PMIs convert that input through mechanisms such as capped profit conversion, escrowed reinvestment, or slice-based disbursement.

2. Internal Redistribution

- Once converted, value is remitted to the HAO and allocated per the ICN's financial architecture:
 - Reinvestment in UMEs
 - Infrastructure enhancement
 - Equity buybacks
 - Reserve fund contributions

This keeps external funding compatible with the network's internal accounting model while offering external participants a familiar interface.

Governance and Oversight

PMIs are legally distinct but bound by a Public Interface Agreement (PIA) with the HAO. Key governance patterns include:

- **Minority External Control:** Majority ownership remains with network entities (e.g., HAO, mature UMEs).
- **Mission Lock:** Legal commitment to the network's stated principles and governance process.
- **Transparency Requirements:** Obligated reporting to HAO councils and auditors.
- **Sunsetting Clauses:** Built-in options for termination, absorption, or spin-out.
- **Interface Stewardship:** Designated individuals or working groups monitor alignment and compliance.

Oversight balances operational independence with accountability to the network.

Risks and Mitigation

PMIs carry risks, including:

Risk	Mitigation Strategy
Capture by external capital	Use structural controls (voting rights, golden shares, sunset clauses).
Cultural dilution	Maintain consistent narrative control, shared branding, and community norms.
Legal overreach	Separate liability clearly, use regulatory firewalls.
Scaling pressure	Cap growth rates or customer acquisition velocity when necessary.

Properly structured, a PMI buffers these risks rather than bypassing them.

Examples of Real-World Analogues

Organization	Type	Relevant Insight
Stocksy United	Member-owned stock platform	Artist-run cooperative interfacing with photo buyers
CoopCycle	License/brand steward	Enforces ethical franchising among local bike co-ops
Zebras Unite Coop	Entrepreneurial coop	Channels VC resistance into long-term cooperative gains
Purpose Foundation	Steward ownership foundation	Locks mission into company charter

Each offers design precedent for network-facing market interfaces.

Relationship to MEE and HAO

- The **MEE** (the network’s protected internal economy) maintains the boundary; PMIs are its economic gateways.
- The **HAO** monitors and governs interface logic; PMIs implement it.
- **UMEs** and **SEPs** (joint ventures between teams) may route market interactions through PMIs to limit external entanglement.

The relationship resembles an embassy: PMIs operate at the edge of the ecosystem, working in both internal protocol and external market terms.

Conclusion

PMIs form the interface between HAO-based networks and the broader economic world, using mission buffering, ownership structuring, and value translation to enable external engagement without requiring changes to internal governance.

With the structural components now complete, the next chapter covers the governance frameworks that hold this system together across scale, time, and difference.

3.0 • Governance

While structure enables coordination, governance ensures alignment, accountability, and evolution over time. In the Humanized Autonomous Organization (HAO) — the network’s coordinating framework — governance is not a fixed hierarchy or a one-time design choice; it evolves in response to complexity, conflict, and change.

This section covers the governance architecture that supports distributed autonomy, consistency with stated principles, and operational resilience. It applies across a network of United Micro Enterprises (UMEs, small, self-managing venture teams of up to ~15 people), Strategic Enterprise Partnerships (SEPs, joint ventures between teams), and ecosystem layers such as the MEE (the network’s protected internal economy).

Instead of centralized control, the HAO combines three governance models:

- **Adaptive Governance Framework (AGF)** (the network’s layered governance system): a polycentric, multi-level structure that distributes decision-making based on subsidiarity and contextual fit.
- **Dynamic Enterprise Agreement (DEA)** (a versioned operating agreement replacing fixed bylaws): defines rules, rights, roles, and protocols across the system.
- Embedded representation and redundancy protocols: reduce single points of failure and bias, and improve responsiveness at every level of the network.

The HAO governance system is designed to be:

- **Participatory:** All members, regardless of position or location, can contribute to and shape their local governance.
- **Responsive:** It adapts to new challenges through versioning, soft-forking, and governance-layer experimentation.
- **Transparent:** Decisions, changes, and governance processes are recorded, published, and made available for network-wide review.
- **Conflict-tolerant:** Conflict is expected and structured into the system rather than avoided or punished. Multiple paths exist for resolution, feedback, and course correction.
- **Technologically supported:** Digital systems (distributed ledgers, role-tracking tools, AI-assisted deliberation platforms) support governance without replacing it.

Within a UME, the MMM ordering — Members, Mission, Market, in that priority sequence (see `docs/appendices/terminology.md` §VII) — sets the order in which a UME weighs the three when they pull against each other. It is an ordering of attention rather than a grant or limit of authority; what a UME may decide without asking anyone is set by the content of its mission, not by MMM. MMM, the ETHICAL Framework, and the PARTS Model are parallel lenses with no general precedence among them; the sole stated exception is that a UME’s mission choice is bounded by ETHICAL values.

This section breaks down the architecture across five dimensions:

- 3.1 Adaptive Governance Framework (AGF)
- 3.2 Dynamic Enterprise Agreement (DEA) Lifecycle
- 3.3 Polycentric Decision Layers
- 3.4 Subsidiarity and Intentional Redundancy
- 3.5 Representation Systems and Feedback Loops

Together, these elements form the governance structure of a HAO-based system, designed to accommodate change while maintaining consistency and broad participation.

§3.1 Membership and Quorum Rules

`docs/03-governance/00-overview.md` introduces the Adaptive Governance Framework (AGF) — the network’s layered governance system — and the Dynamic Enterprise Agreement (DEA) — a versioned operating agreement replacing fixed bylaws — at a structural level, without numeric thresholds. `docs/02-structure/02-united-micro-enterprises.md` states an upper bound on UME (United Micro Enterprise — a small, self-managing venture team, ≤ ~15 people) size only — “up to ~15 people” — with no minimum count and no rule governing how a UME’s membership total interacts with its voting process. This section supplies that missing floor: a minimum-size and vote-parity rule, the bootstrap sequence for the network’s first units, and the size of the founding leadership group. It is the first place in this corpus that states a concrete numeric rule about who may vote, how many members a unit needs, or how votes are weighted.

Source and provenance. The material below is drawn from a single staged extraction of two source documents: `~/code/papr/misc/pio-001-p-context.md` (45 lines, read in full) and `~/code/provide.io/content/english/governance.md` (402 lines, read in full; its governance-relevant content is concentrated in the first ~60 lines). Both are unreviewed working notes — informal founder documents, not ratified specifications — and neither has been checked against a primary or published source. The two agree wherever they overlap. Numbers below are stated as the sources state them; where the two differ in wording rather than value, both wordings are recorded. The source material uses “Cell” or “Cellular Unit” for what this corpus’s current vocabulary calls a UME;

`docs/appendices/lineage.md` §4 records the rename. Quotations below preserve the sources’ original wording; connective text uses UME.

A. Minimum Unit Size and Vote Parity

`pio-001-p-context.md` states the floor and the parity rule in a single sentence: “Each Cellular Unit must have a minimum of 3 member owners, and must always have an odd number member owners.” `governance.md` corroborates the odd-number requirement and adds a stated rationale absent from the first source: “An effort must be made to ensure that each Cell is occupied by an odd number of individuals in order to proactively prevent hung votes.” `governance.md` does not restate the floor of three; the two sources are complementary on this item rather than overlapping.

As stated across both sources, a UME’s membership must therefore satisfy two conditions: at least three member-owners, and an odd total count at every subsequent size, so that a vote among the full membership cannot tie.

This is a composition constraint on total membership, not a turnout requirement. Neither source specifies a minimum-participation threshold for a given vote to be valid — the word “quorum,” in the sense of a minimum share of members who must be present or voting, does not appear in either source, and no equivalent numeric figure does either. What the sources establish is that the pool of eligible voters is kept at an odd size, which prevents ties among however many members participate; it does not state what happens if fewer than the full membership votes.

This composition rule bears on the governance-capture patterns cataloged in `docs/09-adversarial-analysis/02-governance-capture.md` §9.2, particularly the coordination and coalition dynamics described there; that chapter treats the detection and resistance side of the boundary, and its content is not restated here.

B. Bootstrap: Three Seed Units

`pio-001-p-context.md` describes the network’s first units: “To bootstrap the ICN, PIO will form 3 cells, which are intentionally small, and localized, business entities.” The three domains named are finance (managing financial connections, such as between credit unions, micro-loan services, or crowdfunding), property (managing renting, leasing, or purchasing), and service-labor (teams of service-industry workers specializing in areas such as cooking, cleaning, technology, event management, catering, and education).

The source’s own count is not a strict three: it states “3 cells” as the topline figure, then describes the third, service-labor domain as “one, or more” units — allowing that domain to expand into multiple units without stating a maximum. The “3 cells” framing and the “one, or more” qualifier appear in the same passage and are not reconciled there; this section states the ambiguity rather than resolving it to a single number.

C. Founding Leadership Group

`pio-001-p-context.md` states: “provide.io will assemble a leadership team of 5 through 9 entrepreneurial leaders in areas such law, economics, behavioral science, technology, and more.” `governance.md` corroborates the same range with an adjacent but not identical domain list: “The initial cell is composed of 5-9 leaders with expertise in fields such as law, finance, psychology, education, and any other fields that will be required to successfully test provide.io out.”

`governance.md` alone names this group later in the same document: “The ICN is meant to begin with the Officium Cell, which sets up the initial test Cells, and ensures processes, and resource alignment, are constantly being observed by an AI.” The name “Officium Cell” for the five-to-nine-person founding group is sourced only to `governance.md`; `pio-001-p-context.md` does not corroborate it.

§3.2 Voting and Expertise Weighting

`docs/03-governance/00-overview.md` names the Adaptive Governance Framework (AGF) — the network’s layered governance system — as combining “embedded representation and redundancy protocols” with the Dynamic Enterprise Agreement (DEA) — a versioned operating agreement replacing fixed bylaws — but does not describe a ballot mechanism. The only vote-weighting language already in the corpus is at the SEP (Strategic Enterprise Partnership — a joint venture between teams) level: `docs/02-structure/03-strategic-enterprise-partnerships.md` §2.3 lists “**Proportional**: Voting power or influence can be based on stake, effort, or domain expertise” among SEP governance structures. That describes voting between UMEs (United Micro Enterprises — small, self-managing venture teams, ≤ ~15 people) party to a joint venture. This section describes a different relationship — voting among the members of a single UME — and does not restate or extend the SEP-level material.

Source and provenance. As in §3.1, this section is drawn from `~/code/papr/misc/pio-001-p-context.md` and `~/code/provide.io/content/english/governance.md`, both read in full, plus two further documents checked for this topic specifically: `~/code/papr/misc/form.md` and `~/code/papr/bits-form.md`, near-duplicate drafts of the same pitch letter (confirmed by diff to differ in four minor wording points only). All four are unreviewed working notes. `pio-001-p-context.md` carries the only detailed structural description of the mechanism found across the four; `form.md` and `bits-form.md` each contribute a single confirming sentence and no further structural detail.

A. The Hybrid Ballot

`pio-001-p-context.md` states: “provide.io, and the ICN, are democratic entities. They will leverage a combination of one person, one vote, and a ranked/weighted, systems.” The same source continues: “This model will enable domain relevant systems experts to have more influence over a decision that is more relevant to their expertise. Yet allow for someone in a completely different specialization, such as education, to vote on decisions which need to be made which may have more of a systemic impact.”

As stated, the ballot combines two layers: a flat one-person-one-vote baseline, and a ranked or weighted layer that gives subject-matter experts more influence on decisions within their own domain. The flat layer is preserved for decisions the source describes as having “more of a systemic impact” — cross-cutting or

network-wide questions, by contrast with domain-specific ones — where a member outside the relevant domain retains a full vote. The source’s own example contrasts a systems expert’s domain-relevant weight against an education specialist’s vote on a decision of broader, systemic reach.

`governance.md` corroborates that the mechanism applies at the UME level, in a single unelaborated sentence: “The members of the Cell will have weighted voting rights within the provide.io/ICN systems.” It adds no structural detail beyond confirming that weighted voting exists at the unit level. `form.md` and `bits-form.md` — identical at this point in their text — confirm the same pitch point in a sentence with an apparent dropped word: “The ability to vote throughout the organizational network while having a weighted system in place to ensure that subject matter experts are weight[ed more] than someone who has no clue what they’re voting for,” which this section reads as stating that subject-matter experts are weighted more than participants without relevant expertise. Neither document describes the one-person-one-vote half of the hybrid, the domain-versus-systemic distinction, or the learning purpose treated in §3.2.B below — those structural details appear only in `pio-001-p-context.md`.

None of the four sources specifies the mechanics of the weighted layer: no formula for computing a weighted score, no process for certifying who counts as a domain-relevant expert for a given decision, and no rule for combining the flat and weighted layers into a single outcome. This is a gap in the source material, not an omission of this section.

B. Anomalous Votes as a Learning Signal

`pio-001-p-context.md` states, in the sentence immediately following the domain-versus-systemic description quoted above: “This will also enable the SMEs to learn more about anomalous votes.” This is the only statement of this purpose found across the four sources, and it is the entirety of what the source material says about it.

The source frames this as a secondary effect of the weighting mechanism — “also enable” — rather than as its stated primary purpose, which is the influence-weighting described in §3.2.A. As stated, the mechanism’s direction of effect runs from members to experts: a vote that diverges from what subject-matter experts would expect is treated as something for those experts to learn from, not as a vote to be corrected, discounted, or excluded. Nothing in the source frames an anomalous vote as an error to be suppressed or as evidence against the voter; the sentence names only a learning benefit to the SMEs (subject-matter experts) who observe the divergence.

Beyond that single sentence, the mechanics are unstated. The source does not define what counts as “anomalous” — a numeric deviation from an expert-weighted result, a categorical disagreement, or something else. It does not describe what SMEs do with an anomalous vote once observed: whether it triggers a review, a conversation with the voter, a change to how the decision is explained, or nothing beyond passive observation. And it does not state whether the signal is captured or routed through any existing mechanism — for instance, the AGF or Value Alignment Monitoring (VAM) — ongoing checks that actions match stated principles, described in `docs/11-deployment/02-provide-io-as-deployment-engine.md`. No source names either as the channel for this signal; a connection to either would be an addition this section does not make.

This section, together with §3.1, is the first place in this corpus that states a concrete rule for how votes are weighted at the UME level. It supplies the flat and weighted ballot layers and the learning purpose attached to divergence between them; it does not supply the weighting formula, the expertise-certification process, or the anomalous-vote follow-through, because the source material does not supply them either.

§3.3 Accountability Mechanisms

`docs/03-governance/00-overview.md` names two of the five properties it states for the Adaptive Governance Framework (AGF) — the network’s layered governance system — as bearing on accountability without specifying a mechanism for either: “Conflict-tolerant” states that conflict is “structured into the system rather than avoided or punished,” and “Transparent” states that “decisions, changes, and governance processes are recorded, published, and made available for network-wide review.” Neither property names a body that holds a decision-maker to account, a metric that is tracked, or a relationship in which one role answers to another. `docs/07-human-systems/03-conflict-engagement.md` §7.3.3 separately states a five-tier escalation table for conflict between individuals, with a named protocol at each tier, but no trigger condition for moving from one tier to the next, no timeline at any tier, and no mapping of roles to the parties they are accountable to; that gap is not filled here; §7.3.3 remains the corpus’s account of conflict escalation, and this section does not restate it. This section supplies what neither existing passage does: a statement of what is measured, which roles answer to which other roles, and what is stated to follow when a role fails to meet what is expected of it.

Source and provenance. The material below is drawn from a staged extraction of `~/code/mtu-to/`, an unreviewed single-pass output describing the Member Trust Union (MTU) — the network’s credit-union-like financial institution. `docs/13-member-trust-union/00-overview.md` presents the MTU as “a domain-specific implementation of the Humanized Autonomous Organization (HAO) — the network’s coordinating framework” in its opening summary, and states in §13.8 that the MTU is “a HAO instance in its own right” rather than a subordinate layer of one; that framing is this section’s basis for drawing on MTU-specific source material to describe accountability design in the chapter that otherwise treats the network’s governance in general terms. Two source files within the extraction bear on accountability specifically: `mtu-8-5-accountability-mechanisms.md`, read in full for this section, and `mtu-8-6-adaptation-and-evolution-processes.md`, checked and excluded (see the note at the end of this section). The staging extraction describes the material as “deliberately thin” and states that stripping bullet lists naming a category with no attached actor, trigger, or consequence “removed... taxonomy padding” rather than substantive content the writer chose to omit. That thinness carries through to what follows: several of the categories below are named with no mechanism stated for how they operate, and the extraction is explicit that none of the three MTU source files it drew on contains a quantified figure, a stated timeline, or a stated consequence for any accountability relationship it names. Where this section states that the source is silent, that silence is the extraction’s own finding, not a gap introduced in restating it.

A. Accountability Relationships

The source names four accountability relationships directly, each under its own heading in a section titled “Specific Accountability Relationships”: **governance-to-member**, **staff-to-governance**, **MTU-to-community**, and **cross-level** — a local unit’s accountability to the network above it. This is the one place in the source that states a role-to-role accountability pairing rather than a category list, and it is retained here for that reason.

Each pairing in the source carries sub-bullets naming a category with no attached mechanism — for example, “Result Delivery: Achievement of desired outcomes” under one of the four headings — and states no process by which one party in the pair actually holds the other to account: no reporting requirement, no review body, no consequence tied to the relationship. Only the pairing itself is stated here; the category labels beneath each one are not, on the basis that a label naming what accountability is for is not itself a mechanism for how it operates.

B. What Is Measured: Transparency Artifacts

The source’s most concrete content on this topic is a list of named transparency artifacts, under a heading “Transparency Systems”: performance dashboards, decision documentation, financial and resource-use transparency, and impact reporting. The source names each artifact and states no more: no specific indicator or metric populates the “performance dashboards” entry, and no reporting cadence is stated for any of the four.

C. Named Accountability Bodies

Under a heading “Accountability Roles,” the source names five bodies: Ethics Committees, Performance Review Teams, Audit Functions (described as independent), Member Voice Channels, and Supervisory Positions, the last described as staff-facing. The source states no composition for any of the five, no appointment process, and no reporting line — it does not say who sits on an Ethics Committee, who appoints an Audit Function, or to which of the four relationships in §3.3.A any of the five bodies belongs.

D. Peer Accountability

The source names a further channel, described as “team-based” and operating through “structured peer feedback systems,” and distinguishes it explicitly from the hierarchical relationships named in §3.3.A. Beyond the label and the distinction, the source states no mechanics for how peer feedback is structured, how often it occurs, or what is done with it.

E. What Follows

The source states that a “Consequence Framework” exists, and names two components of it: “Progressive Intervention,” described only as “graduated approaches to serious concerns,” and “Pattern Address,” described only as “response to recurring issues.” This is the only place in the source that gestures at a consequence for an accountability failure, and the source is silent beyond the two labels themselves. It does not state what a graduated approach consists of at any step, whether a step carries a defined duration, or what “response” means for a recurring issue — no suspension, removal, restitution, or other named sanction is stated at any tier, and no source states what triggers movement from one component of the framework to the other. This section records that silence rather than supplying a threshold, a deadline, or a penalty the source does not state.

Excluded from this section. `mtu-8-6-adaptation-and-evolution-processes.md` states a four-stage pipeline for revising governance rules — proposal development, review, approval, implementation planning — but names no actor for any stage: no source states who may initiate a proposal, what body reviews it, or what approval threshold applies. This is a rule-amendment process rather than an accountability relationship, and it falls outside what §3.3.A to §3.3.E state; it is not brought in here. The staging extraction separately notes that the source’s only statement of a varying approval threshold — higher thresholds for a fundamental change, more accessible ones for routine refinement — appears solely inside a fabricated attributed quotation in the underlying source, and the extraction drops it rather than presenting it as fact; this section does not reintroduce it.

Boundary with chapter 09. The accountability relationships, bodies, and consequences named above are governance design: who is answerable to whom, and what the source states or leaves silent about what follows. `docs/09-adversarial-analysis/02-governance-capture.md` §9.2.9 and `docs/09-adversarial-analysis/03-framework-capture.md` treat a different question — accountability mechanisms as something a capture pattern can be turned against, stating the case in which “the model’s accountability mechanisms are being asked to act against the people who defined them.” That question, and the detection and resistance mechanisms chapter 09 catalogs against it, are not treated here.

§3.4 Cross-Unit Governance Instruments

`docs/03-governance/00-overview.md` states that the Adaptive Governance Framework (AGF) — the network’s layered governance system — distributes decision-making across United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people — Strategic Enterprise Partnerships (SEPs) — joint ventures between teams — “and ecosystem layers such as the MEE,” and names “Responsive” among the AGF’s stated properties: the system “adapts to new challenges through versioning, soft-forking, and governance-layer experimentation.” Neither passage names an instrument for coordinating between governing units that fall short of forming a SEP, for settling a dispute between units whose jurisdiction overlaps, or for running the versioning and experimentation the “Responsive” property names. `docs/02-structure/03-strategic-enterprise-partnerships.md` and `docs/11-deployment/04-sep-structuring.md` §11.4.3 already specify a SEP Charter — a jointly signed governing document for a new joint-venture entity that UMEs form together, with its own decision protocol, conflict protocols, and exit provisions. This section describes seven further instruments that do not require forming a new entity: a standing coordinating role between units, an agreement format for cross-boundary matters, a bounded space for testing a governance change, a disparity-tracking and transfer proposal, a citizen-audit proposal, a shared contingency form, and a proposed funding form for governance work.

Source and provenance. The material below is drawn from the staged extraction’s `## Countermeasures` section, ultimately sourced to `~/code/papr/polycentric-governance/005-r-claude.md`, `005-r-gemini.md`, and `005-r-chatgpt.md` — three of four model-generated responses to a prompt asking for solutions to the gaps that `docs/09-adversarial-analysis/05-polycentric-failure-modes.md` catalogs as the network’s inherited polycentric failure modes. The extraction’s own assessment, carried forward unchanged in `docs/09-adversarial-analysis/06-countermeasures-and-detection.md` §9.6.20, is that this pass “largely reproduces general public-administration content rather than mechanisms specific to this model,” and that four independent model variants converged on similar answers to a prompt that asked for novel solutions. No claim is made in this section that any named study, author, or research tradition established these seven instruments: the sources are model-generated responses citing no author or paper, on the same footing as the failure-mode material in §9.5 of chapter 09. What follows states what the source proposes, not a finding independently verified against public-administration scholarship.

Each instrument below is one the extraction and `docs/09-adversarial-analysis/06-countermeasures-and-detection.md` §9.6.20 map to specific polycentric failure modes that `docs/09-adversarial-analysis/05-polycentric-failure-modes.md` catalogs by name; those entries are cited by name and section below and are not restated here.

A. Bridging Organizations

The source describes a bridging organization as a standing coordinating role rather than a joint venture: “a network of organizations specifically designed to facilitate coordination/communication/collaboration among governing units, developing shared goals/strategies/metrics and a platform for joint problem-solving and resource-sharing.” As stated, the role is ongoing and cross-cutting — it exists to coordinate between units on a continuing basis, not to jointly own or operate a specific venture, which is what a SEP does. The source states no composition, no appointment process, and no scope limit for a bridging organization; whether one bridging organization serves the whole network or several serve different clusters of units is not addressed.

The AGF already provides a coordinating layer above UMEs and SEPs (`docs/03-governance/00-overview.md`), and `docs/11-deployment/04-sep-structuring.md` §11.4.3 names an existing reviewing role — “the HAO Coordination Node, or an equivalent integrator (e.g., `provide.io`)” — that reviews a SEP Charter for alignment and risk at formation. Neither of those is stated by any source as performing the standing, cross-unit coordinating function described here; a bridging organization, as the source states it, is a distinct institutional role, not a restatement of either.

This instrument is the one the extraction and §9.6.20 map to the coordination cluster that `docs/09-adversarial-analysis/05-polycentric-failure-modes.md` divides between Absence of Overarching Vision (§9.5.1) and Cross-Boundary Problem Paralysis (§9.5.2); those entries are not restated here.

Source: `005-r-claude.md`.

B. Joint Governance Agreements

The source describes a joint governance agreement as a formal agreement between governing units for a cross-boundary matter, without requiring that the units form a new joint entity to hold it: “formal agreements outlining shared responsibilities, commitments, and accountability mechanisms for cross-boundary issues, including joint decision-making, cost-sharing, dispute resolution provisions.” Three components are named directly in that sentence: a joint decision-making provision, a cost-sharing provision, and a dispute-resolution provision. The source states no template for any of the three, no minimum content, and no body that reviews or ratifies an agreement before it takes effect.

A SEP Charter, as `docs/11-deployment/04-sep-structuring.md` §11.4.3 specifies it, already covers cost allocation, decision protocol, and conflict protocols for UMEs that form a joint venture together. The source for this instrument does not state whether a joint governance agreement is meant for units that have not formed a SEP, for a matter narrower than a full joint venture would warrant, or for a relationship between a unit and a layer of the AGF above it rather than between two UMEs; the relationship between the two instruments is not addressed by any source, and this section states that gap rather than resolving it.

This instrument is the one the extraction and §9.6.20 map to the jurisdictional-overlap cluster that `docs/09-adversarial-analysis/05-polycentric-failure-modes.md` catalogs as Conflict Deadlock (§9.5.7), and it bears on Cross-Boundary Problem Paralysis (§9.5.2) through the cross-boundary responsibilities it records; those entries are not restated here.

Source: `005-r-claude.md`.

C. Governance Sandboxes with Sunset Clauses

Two sources name closely related versions of this instrument. `005-r-claude.md` names “policy labs” and “governance sandboxes” as “designated zones where governing units test innovative approaches to complex problems in a controlled setting, piloting new technologies/models/partnerships and generating evidence for broader system change,” and pairs the sandbox form explicitly with “adaptive regulation and sunset clauses,” stated as enabling “time-limited rule testing with built-in monitoring/evaluation/adjustment.” `005-r-gemini.md` independently names “‘Sunset Clauses’ for Policies” as a standalone instrument. As stated across the two sources, the instrument combines a bounded space for testing a governance change with a clause that ends the test automatically unless a further decision extends or adopts it. Neither source states how long a sandbox runs by default, what evidence standard determines whether a tested change is adopted network-wide, or which body authorizes a sandbox in the first place.

`docs/03-governance/00-overview.md` already names “versioning, soft-forking, and governance-layer experimentation” as how the AGF’s “Responsive” property operates, without naming an instrument for it; a governance sandbox with a sunset clause is one way of giving that stated property a concrete form, though no source states that connection, and this section does not adopt it as more than a plausible reading of the two passages together.

This instrument is the one the extraction and §9.6.20 map to the entrenchment and path-dependence cluster that `docs/09-adversarial-analysis/05-polycentric-failure-modes.md` folds into Absence of Overarching Vision (§9.5.1); that entry is not restated here.

Sources: 005-r-claude.md , 005-r-gemini.md .

D. Equity Scorecards and Redistributive Transfers

The source proposes “equity scorecards” that track disparities in outcomes and opportunities across jurisdictions. It states that the resulting information can guide resource allocation and policy priorities, and pairs the scorecards with redistributive mechanisms. It gives inter-jurisdictional transfers and equity funds as examples, intended to make resources and capacities available across communities. The source states neither which outcomes or opportunities count as disparities, how data is collected, nor a threshold, transfer formula, approving body, or review process.

The extraction maps this instrument to the equity and representation cluster that docs/09-adversarial-analysis/05-polycentric-failure-modes.md divides between Elite Capture of Governing Units (§9.5.4) and Equity Drift (§9.5.5). That mapping is the extraction’s; 005-r-claude.md proposes the scorecards and transfers without naming either failure mode.

Source: 005-r-claude.md .

E. Citizen Audits with Open-Data Tracking

The source proposes that citizens play a more active role in monitoring and shaping the performance of a polycentric system through “citizen audits,” participatory budgeting, and other direct-democracy forms. It pairs those forms with digital platforms and open data, stated as making decision-making processes and outcomes trackable across jurisdictions. The source does not specify who is eligible to conduct an audit, what an audit can require, what information may be published, how privacy or retention is handled, or what follows from an audit finding.

The extraction maps this instrument to the accountability and transparency cluster that docs/09-adversarial-analysis/05-polycentric-failure-modes.md catalogs as Accountability Loss Through Opacity (§9.5.8). That mapping is the extraction’s; 005-r-claude.md states the audit and data proposal without naming the failure mode.

Source: 005-r-claude.md .

F. Resilience Networks and Shared Contingency Planning

The source proposes “resilience networks” in which governing units share information, resources, and capacities when shocks and stressors arise. It pairs the networks with shared contingency plans and emergency-response protocols that state roles and responsibilities for different actors in a crisis. The source does not state a trigger for activating a network or protocol, identify an actor, set a resource commitment, or name the body that adopts or authorizes a plan.

docs/06-lifecycle/03-collapse-and-containment-protocols.md already treats containment during a unit’s collapse, and docs/15-advanced/06-systems-simulation-and-stress-testing.md specifies antifragility modeling. The extraction states that resilience networks overlap both areas; it does not say that the proposed network replaces either existing treatment. It maps the proposal to resilience across the failure modes rather than to one named mode.

Source: 005-r-claude.md .

G. Social Impact Bonds for Governance

The source proposes social impact bonds directed at governance initiatives: private investment is channeled to projects addressing gaps in polycentric governance, and success payments to investors depend on predefined governance-improvement metrics. It records the instrument as a funding mechanism rather than assigning it to a particular failure mode. The source does not establish whether the proposed form is compatible with HAO ownership or member-priority arrangements, and it identifies no metric, investor, payer, contract terms, or risk allocation.

Source: 005-r-chatgpt.md .

Instruments not covered here. The same Countermeasures material also maps Knowledge-sharing platforms and adaptive, participatory monitoring and evaluation to chapter 03. Neither is specified here. This section names the two remaining proposals only to delimit its scope; their source mapping appears in docs/09-adversarial-analysis/06-countermeasures-and-detection.md §9.6.20 and its coverage table.

4.0 • Economics

The economic design of the Humanized Autonomous Organization (HAO) — the network’s coordinating framework — governs how capital, revenue, and incentives move through the network. The HAO’s economic architecture directs capital to the network’s value-generating units, the United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people — rather than concentrating it in central administrative layers, as in traditional hierarchical models. This outward allocation is referred to as a trickle-out mechanism. Separately, diminishing contributions return a declining share of revenue to shared infrastructure as UMEs mature and repay their initial allocation.

The model’s governance premise adapts polycentric institutional analysis (Ostrom, 2010) and platform-cooperative ownership debates (Scholz & Schneider, 2016). The capital-allocation mechanics below are HAO design proposals. The framework is designed to:

- **Capital allocation:** Capital enters the network via the HAO and is directed to UMEs based on strategic alignment, demonstrated capacity, and market opportunity. This front-loaded strategy directs capital toward the units nearest to production.
- **Revenue distribution:** A revenue allocation framework returns a higher percentage of generated revenue to the HAO during early stages, funding network infrastructure, repaying initial investments, and supporting UME development. This percentage decreases as UMEs mature, based on predefined milestones, increasing UME autonomy and local reinvestment capacity.
- **Incentive design:** The model specifies separate economic arrangements for individual UMEs and for Strategic Enterprise Partnerships (SEPs) — joint ventures between teams — covering equity distribution, profit sharing, and performance-based compensation at both the individual-enterprise and collaborative-venture level.
- **Reinvestment:** Revenue earmarked for reinvestment is recycled across the network to fund new initiatives, support underperforming UMEs, and maintain shared infrastructure.
- **External capital interface:** For external market interactions, the framework provides a buffer through dedicated public market interfaces, which allow outside investors to engage with the network under terms structured around the trickle-out model rather than standard ROI expectations.

This section describes the components and mechanisms that structure the relationship between capital, labor, and value creation in the HAO model. Investment is directed to the network’s periphery — the UMEs — rather than concentrated centrally.

4.0.1 Devolved Financial Authority and Beyond Budgeting

Beyond Budgeting offers a comparison for the HAO’s proposed allocation rules, not evidence that it shaped the HAO’s original design. In a convenience survey of 201 organizations drawn from Beyond Budgeting Round Table contacts, Matějka, Merchant and O’Grady found that self-reported adopters were more likely than non-adopters to report decentralized decisions, flexible resource allocation and relative targets. The study reports associations in that sample; it does not establish that replacing a master budget would produce the same results in an HAO ([icn-0085](#) in the research bibliography).

In a qualitative genealogy, Becker, Messner and Schäffer describe a tension between presenting Beyond Budgeting as a complete management model and allowing organizations to adapt its component practices ([icn-0072](#)). Hudson’s study reports adoption of individual practices even where the full model was not adopted ([icn-0102](#)). Neither source establishes that incremental adoption is impossible. The HAO’s dynamic allocation arrangements remain design proposals to be specified and tested in their own setting.

Burell and Mattsson’s four-case interview study found varied implementations, with none of the organizations intentionally adopting every Beyond Budgeting principle ([icn-0086](#)). Its respondents also described retained fixed targets and annual expenditure controls in some cases. The authors caution that concurrent organizational changes and the interview design prevent attributing reported outcomes solely to Beyond Budgeting. The cases illustrate implementation variation; they do not validate the HAO’s allocation rules.

§4.1 Trickle-Out Investment Flow

The **Trickle-Out Investment Flow** describes how capital moves through the architecture of the HAO (the network’s coordinating framework). Rather than concentrating in central administrative or executive layers, capital is directed to the network’s edges: directly to value-generating United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people.

4.1.1 Capital Entry and Directionality

In most HAOs, investment capital enters the system via the HAO — a legal and operational entity that coordinates infrastructure, governance, and integration. The HAO does not hold capital as a reserve; it functions as an allocator, distributing capital downstream.

Capital Flow Model:

- **Phase 1 (Seeding):** Capital is directed from the HAO to UMEs selected for strategic relevance, demonstrated capacity, and ecosystem readiness.
- **Phase 2 (Operationalization):** Funds are allocated to support core operating needs, capability development, and go-to-market efforts.
- **Phase 3 (Performance-based Scaling):** The HAO may release additional capital against predefined milestones. Validated-learning cycles (Ries, 2011) can inform that assessment; the capital-release rule is a corpus design proposal.

Capital is distributed unevenly, weighted toward UMEs assessed as most ready to deploy it effectively.

4.1.2 Investment Allocation Criteria

Each capital allocation is guided by a set of dynamic criteria evaluated by either human decision bodies (e.g. governance councils) or a Collaborative Intelligence Network (CIN). Factors include:

- **Strategic Alignment:** How well the UME’s mission and roadmap support the broader network vision
- **Capability Maturity:** Technical, organizational, and operational readiness to deploy capital effectively
- **Market Opportunity Fit:** Near-term and long-term potential for value generation or service delivery
- **Interoperability:** Ability of the UME to collaborate with other nodes — SEPs (joint ventures between teams), other UMEs, or MTUs (the network’s credit-union-like financial institutions)

This approach resembles a cooperative venture studio more than a fixed budgeting process.

4.1.3 Capital Types and Modalities

The HAO framework recognizes several capital types beyond fiat investment:

- **Financial Capital:** Traditional capital deployed as grants, recoverable advances, or equity-aligned allocations
- **Capability Capital:** Investment in upskilling, tooling, and operational infrastructure
- **Reputation Capital:** Access to elevated trust tiers or priority integration into network-wide SEPs
- **Commons Capital:** Non-ownership-based investment in open resources that benefit all UMEs

These capital types can be combined into capital stacks, depending on a UME’s lifecycle stage or strategic function.

4.1.4 Diminishing Contribution Protocol

The **diminishing contribution protocol** governs how value returns to shared HAO infrastructure: as a UME matures and repays its initial capital — whether measured financially or in ecosystem contribution — the required contribution decreases over time.

UME Stage	% Revenue Returned to HAO	Primary Purpose of Return
Seeding	30–40%	Infrastructure funding, repayment, support of new UMEs
Early-Stage	20–25%	Network maintenance, scaling support
Mature	10–15%	Strategic reinvestment and reserves

Under this model, early investments are repaid and capital continues to circulate through the network, while UMEs gain more financial autonomy as they stabilize.

4.1.5 Temporal and Cyclical Flow Design

To support stability and coordination, the framework applies multiple time-based distribution cycles:

- **Monthly:** Operational distributions for baseline function and liquidity
- **Quarterly:** Performance-based recalibration or growth fund participation
- **Annual:** Strategic redistributions to fund long-term infrastructure, research, or expansion initiatives

This rhythm is intended to keep flows predictable while allowing adjustment during market or ecosystem shifts.

4.1.6 Flow Transparency and Ledgering

All investment flows are recorded by default in the **Distributed Ledger Infrastructure (DLI)**. This enables:

- Historical tracking of capital deployment
- Auditable records of performance and return
- Public visibility (if permitted) into the flow of commons-based capital
- Interoperability with trust verification mechanisms

The ledger also supports automated threshold-based adjustments to HAO contribution percentages, triggered by predefined financial or performance milestones.

4.1.7 Comparative Context

The trickle-out model differs from:

- **Corporate Venture Capital**, where value tends to accumulate at the holding level
- **Top-down NGOs**, which centralize allocation decisions in a governing body rather than distributing them to production units
- **Blockchain-based DAOs**, which typically lack stage-based capital protocols

The HAO model combines elements of venture financing, cooperative principles, and reserve-based resilience design.

4.1.8 Summary

The Trickle-Out Investment Flow structures the HAO's economics around directing capital to the network's edges, adapting flows to UME lifecycle phases, and reducing central claims over time as UMEs mature.

Key References:

- Ostrom, E. (2010). *Beyond markets and states: Polycentric governance of complex economic systems*.
- Scholz, T., & Schneider, N. (Eds.). (2016). *Ours to Hack and to Own: The Rise of Platform Cooperativism, a New Vision for the Future of Work and a Fairer Internet*. OR Books.
- Ries, E. (2011). *The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*.

§4.2 Revenue Allocation Framework

The **Revenue Allocation Framework (RAF)** defines how value generated within a Humanized Autonomous Organization (HAO) — the network's coordinating framework — is distributed across the network. It builds upon the trickle-out investment flow by specifying how, when, and to whom revenue flows once UMEs (small, self-managing venture teams) and SEPs (joint ventures between teams) begin generating economic returns.

This framework is designed to (1) support the sustainability of network infrastructure, (2) reward contributors according to defined criteria, and (3) maintain incentives for local autonomy and collaborative behavior. The framework distributes revenue based on contribution and resilience rather than maximizing profit.

4.2.1 Allocation Tiers and Distribution Logic

Revenue generated by a UME or SEP is allocated into distinct **tiers**, each serving a specific systemic purpose. These tiers may be applied sequentially or in parallel, depending on the implementation.

Standard Allocation Tiers:

1. **Operating Reserve Tier** – Maintains a liquidity reserve for each UME or SEP (e.g. 3–6 months of baseline expenses).
2. **Base Compensation Tier** – Pays team members their guaranteed minimum or cooperative base wage.
3. **Network Contribution Tier** – Allocates a percentage to the HAO (as defined by the diminishing contribution protocol).
4. **Profit Sharing Tier** – Distributes surplus to team members and contributors based on internal agreements.
5. **Reinvestment Tier** – Channels a portion of profits into UME-level or network-wide funds (e.g. innovation, training, mutual aid).
6. **Commons Maintenance Tier** – Optional: contributes to shared resource pools, licensing upkeep, or ecological commons if relevant.

These tiers can be parameterized by governance vote, automated by smart contracts, or manually managed using dynamic accounting systems.

4.2.2 Temporal Dimensions of Distribution

The RAF structures revenue distribution into cycles:

- **Weekly or Biweekly:** Base compensation payouts for members
- **Monthly:** Network contribution reconciliation and reserve top-ups
- **Quarterly:** Performance-based distributions and reinvestment fund replenishment
- **Annual:** Strategic redistributions, surplus allocations, or special dividends

This timing is intended to avoid early depletion of returns and support predictable financial planning.

4.2.3 Contribution Mapping

HAOs use contribution mapping frameworks to determine how profit-sharing and performance-based compensation is calculated.

Methods include:

- **Slicing Pie-style Dynamic Equity Models** (Moyer, 2012): Allocate equity based on risk-adjusted time and resource contributions.
- **Reputation-weighted Distributions:** Uses trust and verification scores to modulate shares.
- **Role- or Task-based Shares:** Aligns with functional responsibility and deliverables.

These can be hybridized to support different UME cultures, lifecycle stages, or mission types.

4.2.4 HAO Contribution Adjustment Protocol

The HAO's share of revenue adjusts based on:

- UME lifecycle maturity (see Section 4.1.4)
- Risk exposure of the HAO (e.g. investment amount, legal shielding, shared infrastructure)
- Current health of network-wide systems (e.g. underfunded maintenance funds trigger contribution spikes)

This behavior is modeled algorithmically, via multi-party governance, or both. The goal is to balance network needs against local autonomy.

4.2.5 Strategic Enterprise Partnership (SEP) Revenue Logic

SEPs operate as economic bridges between UMEs, or between a UME and external interfaces. Their revenue logic differs slightly:

- Revenue splits are determined at the time of SEP formation, based on contributed resources and projected value.
- SEP income is routed back to participating UMEs using agreement-specific formulas.
- A percentage is optionally directed to a **SEP Reinvestment Pool**, governed jointly by the partners.
- The HAO receives a small SEP-specific contribution (typically 3–10%), unless it played a material role in co-founding or resourcing the SEP.

This structure supports shared risk and shared return between SEP partners.

4.2.6 Buffering and Reserve Strategies

Revenue distributions are buffered by reserve protocols, designed to:

- Smooth volatility across business cycles
- Enable proactive infrastructure investment
- Maintain liquidity during downturns

UMEs are encouraged (or required) to maintain minimum operating reserves, while the HAO and SEPs maintain network-wide or shared reserves. These reserves function as shock absorbers and working capital pools.

4.2.7 Alignment with Network Principles

All allocations must be traceable and aligned with the **ETHICAL** and **PARTS** frameworks (Section 2.5):

- **Transparency:** Every distribution is ledgered and visible (internally, externally, or both)
- **Collaboration:** Distribution rules are co-designed and evolve via governance
- **Resilience:** No single point of failure or dependency in the flow chain
- **Longevity:** Profit is recycled and reinvested rather than depleted at each cycle

4.2.8 Summary

The Revenue Allocation Framework formalizes how revenue flows within the HAO network, across members, the HAO, reinvestment vehicles, and collaborative partnerships. It is designed to scale with network growth, reinforce local autonomy, and avoid centralizing financial power.

Key References:

- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.
- Moyer, M. (2012). *Slicing Pie: Funding Your Business Without Funds*. Lake Shark Ventures.
- Raworth, K. (2017). *Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist*. Random House Business Books.
- Scholz, T., & Schneider, N. (Eds.). (2016). *Ours to Hack and to Own: The Rise of Platform Cooperativism, a New Vision for the Future of Work and a Fairer Internet*. OR Books.

§4.3 Strategic Enterprise Partnership (SEP) Economic Agreements

Strategic Enterprise Partnerships (SEPs) — joint ventures between teams — are structured collaborations between two or more United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people — and, occasionally, the HAO (the network's coordinating framework), formed to pursue shared goals, market opportunities, or infrastructure development. While UMEs operate with significant autonomy, SEPs provide a formalized interface for joint value creation. This section defines how economic agreements within SEPs are structured, negotiated, and executed over time, with attention to fairness, accountability, and alignment with the HAO's trickle-out economic principles.

4.3.1 SEP Formation and Agreement Principles

SEPs are governed by explicit **collaboration charters** or **SEP Economic Agreements (SEAs)**. These agreements typically define:

- **Purpose and Scope:** Clear articulation of the joint initiative's goals and deliverables
- **Participants and Roles:** List of contributing UMEs (or HAO) and their respective roles
- **Resource Commitments:** Inputs provided by each participant—capital, labor, IP, infrastructure
- **Value Sharing Logic:** How output (revenue, assets, equity, commons) will be distributed
- **Governance Mechanisms:** Dispute resolution, consent thresholds, amendment procedures
- **Lifecycle and Exit Conditions:** Triggers for ending or evolving the partnership

The SEA acts as both a legal and operational contract, with optional smart contract enforcement if supported by the underlying platform.

4.3.2 Value Distribution Models

SEP revenue and value distribution varies depending on the nature of the collaboration. Common patterns include:

- **Proportional Contribution Model:** Distributions are allocated according to initial or ongoing contributions (e.g., a 40/30/30 split based on resources committed).
- **Equity-Pool Hybrid:** A fixed portion of SEP revenue goes into a shared equity pool or tokenized representation, with periodic profit-sharing.
- **Outcome-Based Model:** Distribution is based on performance metrics, such as customer acquisition, service delivery, or uptime.
- **Commons-Contribution Allocation:** If the SEP creates commons (e.g., open-source infrastructure, knowledge), value is allocated toward maintenance and governance first.

These models are not mutually exclusive and are often hybridized to reflect the complexity of multilateral ventures.

4.3.3 SEP Capitalization and Reinvestment

SEPs may be:

- **Self-funded** by participating UMEs
- **Cofinanced** by the HAO (particularly for infrastructure-level SEPs)
- **Externally capitalized** via Public Market Interfaces (PMIs) — buffer companies between the network and outside investors — or mission-aligned investors

Regardless of source, SEPs are encouraged to maintain:

- **Operating Reserves:** For liquidity and buffer capacity
- **Reinvestment Funds:** Directed toward SEP infrastructure or second-generation collaborations
- **Redundancy Funds:** Used to address operational failure, team turnover, or regulatory shifts

Reinvestment decisions are made via multi-party governance, typically with each participating UME or stakeholder holding weighted decision rights.

4.3.4 SEP Lifecycle Revenue Flow

SEP economics evolve over time. A standard lifecycle might include:

- **Stage 1 – Seeding:** Founding UMEs receive more favorable revenue shares to reward risk.
- **Stage 2 – Operational Maturity:** Revenue stabilizes, and profit-sharing becomes standardized.
- **Stage 3 – Replication or Forking:** SEPs may spin off into new UMEs or license their model to other networks, generating new revenue flows (royalties, licenses, fractional ownership).

This lifecycle-based flow is designed to reward early contributors while maintaining adaptability for future evolution.

4.3.5 Role of the HAO in SEPs

The HAO may serve multiple functions in a SEP:

- **Neutral Arbiter:** Enforces baseline ethics, trust frameworks, and dispute mediation
- **Co-contributor:** Offers infrastructure (e.g. legal templates, digital identity systems)
- **Strategic Investor:** Provides seed capital, coordination, or platform exposure
- **Governance Steward:** Participates in meta-governance if the SEP spans multiple domains

In such cases, the HAO may receive a small share (typically 3–10%) of revenue for its enabling role, which diminishes over time unless renewed by stakeholders.

4.3.6 SEP Exit and Succession Planning

Because SEPs are semi-autonomous, exit protocols are codified upfront:

- **Dissolution Clauses:** Define what happens to unspent capital, shared IP, and revenue pipelines
- **Continuation Options:** Remaining UMEs may elect to continue under a revised agreement
- **Spinout Mechanisms:** High-performing SEPs may be forked into new UMEs or institutionalized within the ICN (the reference cooperative business network) or MTU (the network's credit-union-like financial institution) as repeatable patterns

This clarity reduces conflict and preserves continuity of value through the transition.

4.3.7 Commons-Based SEPs

Some SEPs are formed not to generate profit, but to produce shared infrastructure or cultural goods. These are called **Commons-SEPs**, and they:

- Receive grant-based or pooled funding
- Generate non-rival assets (e.g. open standards, APIs, educational media)
- Are governed by guardianship trusts or multistakeholder boards
- Return value via enablement (e.g. enabling multiple UMEs to generate profit)

Revenue, if any, is recycled into maintenance or ecosystem-wide investment. Commons-SEPs support capacity-building and alignment across the network.

4.3.8 Summary

Strategic Enterprise Partnerships extend the economic logic of HAOs from single-enterprise to multilateral collaboration. Their agreements function as economic architectures, defining how risk, reward, and control are shared. SEPs combine cooperative principles, flexible financing, and reinvestment to support economic scaling, innovation, and culture propagation across the HAO ecosystem.

Key References:

- Scholz, T. (2016). *Platform Cooperativism: Challenging the Corporate Sharing Economy*. Rosa Luxemburg Stiftung.
- Ostrom, E. (2010). *Beyond Markets and States: Polycentric Governance of Complex Economic Systems*.
- Kelly, M. (2012). *Owning Our Future: The Emerging Ownership Revolution: Journeys to a Generative Economy*. Berrett-Koehler.
- Bollier, D., & Helfrich, S. (Eds.). (2015). *Patterns of Commoning*. Commons Strategies Group.

§4.4 Equity and Member Compensation Models

A defining feature of the Humanized Autonomous Organization (HAO) — the network's coordinating framework — is its approach to aligning incentives across the network without concentrating ownership returns in a small shareholder class. Equity and compensation models within the HAO are designed to distribute value equitably, reward contributions transparently, and support long-term commitment to the system's mission, without concentrating power or wealth.

This section outlines the principles, mechanisms, and lifecycle dynamics of member compensation across UMEs (small, self-managing venture teams, $\leq \sim 15$ people), SEPs (joint ventures between teams), and the HAO itself, highlighting both monetary and non-monetary forms of equity and benefit.

4.4.1 Guiding Principles of Compensation and Ownership

The equity and compensation system within HAOs is structured around the following foundational principles:

- **Alignment over Accumulation:** Ownership reflects contribution and stewardship, not passive capital gains.
- **Multi-form Equity:** Value is stored in financial shares as well as in trust, access, and decision rights.
- **Temporal Fairness:** Early contributors are rewarded, but mechanisms prevent permanent asymmetry.
- **Portable Recognition:** Contributor status and performance history can travel across UMEs and SEPs.

These principles are the corpus's synthesis. They draw on dynamic, contribution-proportional equity (Moyer, 2012) and the broader generative-ownership tradition (Kelly, 2012), but the four-part formulation is not attributed to either source.

4.4.2 Compensation Modalities

Compensation within HAOs includes a mixture of:

Type	Definition	Examples
Base Compensation	Regular payment for labor, usually at or above local living wage	Weekly/biweekly payments
Performance Distribution	Additional income based on UME/SEP profitability and member contribution	Quarterly bonuses, surplus distributions
Profit Share / Dividends	Percentage of retained earnings or net surplus shared based on role or share	Annual dividend
Equity Accrual	Ownership stake in the UME/SEP reflected in internal accounting or token	Slices, reputation-weighted shares

Type	Definition	Examples
Network Profit Pooling	Allocation from overall ICN (the reference cooperative business network) or MTU (the network's credit-union-like financial institution) growth tied to long-term engagement	Cooperative patronage-style payouts
Non-Financial Value	Access to services, decision rights, healthcare, reputation, or housing	Community services, voting shares, housing

This multi-dimensional approach supports both near-term liquidity and long-term ownership, without vesting structures that forfeit accrued equity on departure.

4.4.3 Dynamic Equity Allocation (Slicing Model)

UMEs and SEPs may adopt a **Dynamic Equity Model**, adapted from the *Slicing Pie* framework (Moyer, 2012), where:

A member's equity share = Their proportion of what they put at risk.

This includes time, money, intellectual property, and other resources. Key characteristics:

- **Normalized Slices:** Time and cash contributions are weighted differently (e.g., 1 cash unit = 4 slices, 1 time unit = 2 slices)
- **Rolling Adjustments:** Equity shares update as new contributions occur
- **Departures and Recoveries:** When a member leaves, their slice may convert to a fixed percentage or be bought out
- **Trust-Gated Recovery Tiers:** As a corpus extension, behavior at exit may influence final equity recovery rights, aligned with MTU trust layers

This model prevents equity freezes and aligns rewards with real-time contributions and risk exposure.

4.4.4 SEP-Specific Equity Models

In Strategic Enterprise Partnerships (SEPs), equity is allocated by a joint charter and may take the form of:

- **Joint Contribution Shares:** Proportional to input (labor, IP, capital)
- **Performance-weighted Pools:** Based on milestone achievement
- **Deferred Stake Mechanisms:** Equity vests after defined collaboration duration
- **SEP Commons Pools:** Where part of the value is intentionally non-owned and shared

Equity in SEPs may be redeemable, transferable, or set to decay over time, depending on the mission, duration, and agreement terms.

4.4.5 HAO-Level Participation and Long-Term Accrual

Participants in long-running HAOs (e.g., through governance, R&D, incubation roles) may receive:

- **Ecosystem Equity Shares:** Similar to cooperative patronage dividends
- **Network-Wide Profit Shares:** Based on reputation, longevity, and validated contribution
- **Time-Banked Ownership:** A model where long-term coordination hours translate into HAO-equity (fungible with future influence or income)

These mechanisms are intended to build institutional memory, member loyalty, and continuity, while limiting the concentration of long-term influence among a small group.

4.4.6 Vesting, Lockups, and Exit

To prevent speculation or premature withdrawal of value:

- **Lockup Periods:** Equity becomes redeemable over time (6–36 months)
- **Purpose-Tied Vesting:** Shares are contingent on fulfilling impact, mission, or contribution commitments
- **Exit Recovery Frameworks:** Based on behavior, departures are categorized (e.g. fair, neutral, toxic), with equity scaled accordingly

This mechanism is designed to support exits without prolonged institutional disputes.

4.4.7 Transparency and Governance Integration

All equity and compensation models are:

- **Fully Transparent** (internally, at minimum): Members can view slices, shares, performance tiers
- **Governed Collectively:** Changes require governance approval

- **Recorded on Distributed Ledger Infrastructure (DLI):** Enabling auditing, integrity, and cross-UME traceability

These integrations extend equity beyond a financial instrument to also reflect community trust and contribution.

4.4.8 Summary

The HAO's compensation and equity model decouples ownership from capital alone, tying it instead to risk, contribution, trust, and long-term commitment. It is designed to support member sufficiency and distributed decision-making power while limiting the centralization of ownership and control.

This model aims to support:

- Motivation without a fixed hierarchy
- Fairness without rigid rules
- Growth without concentrating returns in a small ownership class

Key References:

- Moyer, M. (2012). *Slicing Pie: Funding Your Business Without Funds*. Lake Shark Ventures.
 - Ostrom, E. (2010). *Beyond Markets and States: Polycentric Governance of Complex Economic Systems*.
 - Kelly, M. (2012). *Owning Our Future: The Emerging Ownership Revolution: Journeys to a Generative Economy*. Berrett-Koehler.
 - Raworth, K. (2017). *Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist*. Random House Business Books.
-

§4.5 Reinvestment Mechanisms

Traditional economic systems typically prioritize returns to shareholders or centralized capital accumulation. The HAO (the network's coordinating framework) instead treats reinvestment as a structural default, keeping capital in circulation to support future contributors, expand capacity, and build resilience.

Reinvestment mechanisms are deployed at three levels:

1. UME-Level Reinvestment
2. SEP-Level Collaborative Reinvestment
3. HAO-Level Strategic Reinvestment

Each addresses a different timescale and scope of impact, following the design principles described in §4.5.1.

4.5.1 Principles of Reinvestment in the HAO

The reinvestment model is governed by the following design logics:

- **Reinvestment Flow:** Surplus is repurposed to fund future opportunity.
- **Embedded Solidarity:** Reinvestment supports struggling, emerging, or high-impact initiatives.
- **Decentralized Design:** Entities self-manage reinvestment via formal or autonomous protocols.
- **Dynamic Allocation:** Reinvestment ratios are adjustable based on lifecycle stage, network health, or local priorities.

This is intended to create an adaptive economic loop in which growth is recursive rather than strictly linear.

4.5.2 UME-Level Reinvestment

UMEs (small, self-managing venture teams, $\leq \sim 15$ people) are required, or strongly encouraged, to maintain internal reinvestment pools funded from their operating surplus. These pools are used for:

- **Infrastructure Expansion** (e.g., tooling, facilities)
- **Workforce Upskilling**
- **Resilience Buffers** (e.g., savings for downturns)
- **Ecosystem Projects** (e.g., contributing to commons or funding adjacent UMEs)

Typical baseline: 10–20% of net surplus is allocated to this pool monthly or quarterly.

Governance over these funds resides with members, not founders or executives, consistent with the ICN (the reference cooperative business network)'s principle of worker-aligned capital control.

4.5.3 SEP Reinvestment Pools

Strategic Enterprise Partnerships (SEPs) — joint ventures between teams — often include joint reinvestment clauses in their Economic Agreements (see §4.3), such as:

- **Percentage of Net Revenue Reallocated** to a shared innovation fund
- **Royalties on Derived Products** directed into commons expansion
- **Milestone-Triggered Capital Pools** for follow-on phases or adjacent venture incubation

These funds are governed by joint stewardship councils or multi-signature smart contracts to ensure no single UME dominates decision-making.

Example: A logistics SEP operating across three UMEs allocates 15% of quarterly profit into an R&D fund to improve supply chain analytics for the entire network.

4.5.4 HAO-Level Strategic Reinvestment

The HAO maintains network-wide reinvestment mechanisms sourced from:

- Contributions from UMEs and SEPs (see §4.2)
- Returns from external market interfaces (e.g., Contribulo or CoopCycle models)
- Philanthropic or aligned institutional funding
- Long-term surplus from HAO-coordinated ventures

These funds are deployed to:

- **Seed New UMEs**
- **Subsidize Underperforming Units** with high mission alignment
- **Invest in Infrastructure** (e.g., ledger systems, MTU (the network's credit-union-like financial institution) expansion, AI models)
- **Support Ecosystem Innovation**, particularly cultural or experimental initiatives

A portion of HAO reinvestment is allocated via **participatory budgeting**, allowing members to vote on priority projects.

4.5.5 Reinvestment Protocol Design

Each reinvestment mechanism is defined by:

- **Trigger Events:** Surplus thresholds, milestone completions, or ecosystem alerts
- **Allocation Ratios:** Formula-based, percentage-based, or voted distributions
- **Time Horizon:** Immediate (1–3 months), mid-term (1–2 years), or long-term (5+ years)
- **Governance Mechanism:** Local vote, network quorum, or delegated domain control

To support resilience under stress, many protocols include redundancy planning, such as:

- Overlapping funds for mission-critical functions
 - Rotating steward councils
 - Reinvestment decay timers (use-it-or-lose-it clauses)
-

4.5.6 Commons and Mission-Aligned Reinvestment

Some reinvestment explicitly supports non-monetizable value creation, such as:

- Open-source technology contributions
- Educational infrastructure
- Art, culture, or wellness systems
- Ecological restoration or land stewardship

This aligns with the **Commons-SEPs** defined earlier, and with HAO principles of intergenerational stewardship and member well-being.

4.5.7 Multi-Capital Reinvestment

Not all reinvestment is financial. HAOs may reinvest:

- **Social Capital** (trust, reputation access, introductions)
- **Technological Capital** (shared platforms, dev time)
- **Cultural Capital** (knowledge, language, rituals)
- **Natural Capital** (land, ecosystems, energy)

This multi-capital approach to reinvestment is intended to support resilience across domains and reduce dependence on any single form of capital.

4.5.8 Transparency and Auditing

All reinvestment flows—financial or otherwise—are:

- **Ledgered in the DLI**
- Subject to periodic audits (internal or third-party)
- Evaluated against impact metrics aligned with network values

Optional tools include:

- **Impact Dashboards**
- **Reinvestment Maps** (graphing flows over time)
- **Participatory Reinvestment Simulations**

These tools are intended to increase visibility, build trust, and support accountability.

4.5.9 Summary

Reinvestment in the HAO is treated as a systemic behavior rather than an occasional decision, built into the design of the ICN. The intent is for value to circulate forward, for capacity to compound over time, and for participants to be supported across multiple cycles of initiatives.

Through decentralized, multi-layered reinvestment protocols, the HAO aims to support economic sustainability and long-term collective capacity.

Key References:

- Raworth, K. (2017). *Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist*. Random House Business Books.
 - Ostrom, E. (2010). *Beyond Markets and States: Polycentric Governance of Complex Economic Systems*.
 - Kelly, M. (2012). *Owning Our Future: The Emerging Ownership Revolution: Journeys to a Generative Economy*. Berrett-Koehler.
 - Fairbairn, B. (2003). *Three Strategic Concepts for the Guidance of Co-operatives: Linkage, Transparency, and Cognition*. University of Saskatchewan, Centre for the Study of Co-operatives.
-

§4.6 External Interface Economics and Investor Buffers

The HAO (the network’s coordinating framework) model treats full economic autonomy as compatible with selective, structured engagement with external financial systems, including traditional markets, institutional investors, and regulatory environments. Rather than excluding external capital, the model defines **Investor Buffer Interfaces (IBIs)** — structured, semi-permeable interfaces that translate between external investment logic and internal trickle-out economics.

These mechanisms protect UMEs (small, self-managing venture teams, $\leq \sim 15$ people) and internal actors from market volatility, pressure to prioritize short-term returns over network goals, and misaligned incentives, while still enabling capital inflow, liquidity access, and regulatory bridging where beneficial.

4.6.1 Purpose of External Interfaces

External interfaces serve three functions:

1. **Capital Translation** – Converting external investment into capital that is productive within the HAO and does not draw disproportionate returns out of the network.
2. **Risk Insulation** – Shielding local economic actors (UMEs, MTUs (the network’s credit-union-like financial institution), and SEPs (joint ventures between teams)) from the distortions of speculative or short-term financial incentives.
3. **Narrative Bridging** – Helping external stakeholders understand and interact with the HAO’s different economic logic through familiar frames (e.g., equity, returns, governance).

These interfaces combine financial function with embedded values, trust mechanisms, and dynamic constraints.

4.6.2 Types of External Interface Entities

The ICN (the reference cooperative business network) can deploy multiple types of external interface structures, including:

Interface Type	Function	Example
Public Market Interface (PMI) — a buffer company between the network and outside investors	Bridges HAO entities to public investors via equity-like instruments	Contribulo
License-Based SEP	Offers commercial licensing of commons-developed tech	CoopCycle's license model
Joint Ventures with Traditional Firms	Co-creates market-facing products with safeguards	SEP w/ limited external capital
Federated Holding Trust	Aggregates partial UME ownership under cooperative governance	Cooperative Investment Funds
Tokenized Ecosystem Access	Offers time-bound, scoped token access to network assets	Utility-token gated APIs

Each is tailored to context, legal jurisdiction, risk profile, and mission alignment.

4.6.3 Ownership and Governance Protections

To prevent mission drift and economic enclosure, all external interfaces are bound by:

- **Majority HAO/UME Ownership:** The allocation for the Contribulo example remains pending an owner decision
- **Golden Governance Shares:** Non-transferrable veto or override rights held by the HAO or MTU
- **Purpose-Locked Articles:** Corporate charters that legally prevent deviation from foundational principles
- **Time-Bound Concessions:** Investor rights sunset after defined ROI or repayment period

These provisions are intended to keep the external interface subordinate to the network, rather than the reverse.

4.6.4 Flow Control and Capital Translation

External investment entering through these interfaces is not transferred directly to UMEs or SEPs. It passes through **translation protocols**, such as:

1. **Capital Conditioning:** Funds are deployed as recoverable grants, milestone-based tranches, or capped-revenue-sharing instruments, not traditional equity
2. **Purpose Conversion:** Investment is earmarked for infrastructure, commons production, or capability-building, not profit distribution
3. **Flow Dampening:** Internal entities receive capital over time, reducing boom-bust behavior and speculative pressure

4.6.5 Return Structures for External Investors

Investor returns follow the structures below:

- **Capped ROI:** Fixed-multiple or time-bound return expectations (e.g., 2x return within 5 years)
- **Revenue Share Agreements:** Tied to specific products or external-facing services
- **Tokenized Dividends:** Non-voting tokens entitling holders to a slice of interface-specific revenue
- **Exit Through Use:** Investors receive access, participation rights, or licensing, rather than liquidation

The goal is finite engagement aligned with impact outcomes, rather than a perpetual return stream.

4.6.6 Selective Permeability & Market Firewalls

External interfaces maintain **firewall policies**, such as:

- No direct investor exposure to UME operations
- No equity stakes in commons or internal governance
- No claims on member-level compensation or profit pools
- No influence on core protocol evolution

This preserves the internal cultural coherence and autonomy of the ICN.

4.6.7 Trust and Transparency in External Interfaces

Investor interfaces must earn and maintain trust through:

- **Immutable Commitments** (e.g., DLI-recorded charters, trust contracts)
- **Public Impact Reporting** (aligned with ETHICAL framework)
- **Third-Party Oversight** (via MTUs, cooperative federations, or rotating steward councils)

These mechanisms are intended to enforce alignment through social and legal accountability rather than market dynamics.

4.6.8 Use Case: Contribulo (Illustrative)

In the Contribulo model:

- **Ownership:** The allocation between network stewardship and public investors remains pending an owner decision
- **Function:** Aggregates external revenue from logistics tools and service APIs licensed to traditional market players
- **Capital Use:** Reinvested into ICN-wide infrastructure (e.g., MEE (the network's protected internal economy) provisioning, AI co-pilots)
- **Return Path:** Public investors receive capped dividends from licensing revenues, rather than direct equity in UMEs

This illustrates one way profit aligned with HAO principles can be generated without altering internal equity or autonomy arrangements.

4.6.9 Summary

The External Interface and Investor Buffer system allows the HAO to engage with outside capital while maintaining economic and cultural autonomy. These interfaces are designed to:

- Absorb aligned capital
- Translate and buffer investment flows
- Protect contributors and commons
- Maintain systemic integrity

Taken together, they are intended to provide access to external resources without altering the core structure of the ICN.

Key References:

- Kelly, M. (2012). *Owning Our Future: The Emerging Ownership Revolution: Journeys to a Generative Economy*. Berrett-Koehler.
- Ostrom, E. (2010). *Beyond Markets and States: Polycentric Governance of Complex Economic Systems*.
- Bollier, D., & Conaty, P. (2016). *Democratic Money and Capital for the Commons: Strategies for Transforming Neoliberal Finance Through Commons-Based Alternatives*. Commons Strategies Group and Heinrich Böll Foundation.
- Scholz, T. (2016). *Platform Cooperativism: Challenging the Corporate Sharing Economy*. Rosa Luxemburg Stiftung.
- Raworth, K. (2017). *Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist*. Random House Business Books.

6.0 • Lifecycle

Introduction

Lifecycle management within a Humanized Autonomous Organization (HAO) — the network’s coordinating framework — governs the progression, adaptation, and eventual transition or dissolution of its component entities, particularly **United Micro Enterprises (UMEs)** — small, self-managing venture teams of up to ~15 people — and **Strategic Enterprise Partnerships (SEPs)** — joint ventures between teams — while maintaining coherence with the broader system’s values, integrity, and operational resilience.

Unlike conventional organizations, where lifecycle stages are dictated by financial metrics or external market pressures, the HAO framework grounds lifecycle transitions in principle-aligned criteria, socio-technical health, and context-aware governance protocols. This section outlines the models and mechanisms through which entities within the HAO ecosystem emerge, evolve, collaborate, degrade, or conclude their participation.

A key distinction in HAO lifecycle design is the absence of rigid top-down control. Instead, entities self-organize and evolve in alignment with shared values encoded in the **Dynamic Enterprise Agreement (DEA)** — a versioned operating agreement replacing fixed bylaws. This is intended to support both bounded autonomy and systemic coherence, allowing the network to remain adaptive and to strengthen under stress rather than merely resist it.

Lifecycle protocols are built around four primary needs:

- 1. **Genesis** – Enabling the intentional creation of new UMEs or SEPs, informed by network needs, member capabilities, and strategic alignment.
- 2. **Evolution** – Supporting capacity development, maturity transitions, and reconfiguration in response to environmental or internal changes.
- 3. **Integration & Exit** – Managing collaboration across UMEs or with external entities (e.g., via SEPs), and defining ethical offboarding mechanisms when dissolution becomes necessary.
- 4. **System Regeneration** – Facilitating learning, reinvestment, and adaptive changes to governance and operational models based on lifecycle feedback.

These stages are accompanied by monitoring systems (e.g., Value Alignment Monitoring, Enterprise Culture Cultivation), decision-support protocols, and institutional memory tools (e.g., collaborative intelligence systems), so that transitions are intentional, transparent, and recoverable.

Lifecycle management is also a strategic dimension of network health, not solely a matter of operational continuity. As in ecological systems, decay and renewal play a role in resilience. A well-structured HAO lifecycle framework is intended to reduce the chance that an entity persists past its usefulness, that failures become systemic, or that lessons are lost over time.

This section introduces formal lifecycle models, stage-specific governance triggers, and process templates that guide HAO constituents through their lifecycles, with the aim of keeping autonomy and accountability in productive tension throughout the system’s evolution.

\$6.1 UME Lifecycle and Transition States

The lifecycle of a **United Micro Enterprise (UME)** — a small, self-managing venture team ($\leq \sim 15$ people) — is a structured, adaptable process that supports the emergence, evolution, maturation, and potential dissolution of value-producing entities within a Humanized Autonomous Organization (HAO) — the network’s coordinating framework. Unlike conventional startups or business units, UMEs operate under conditions of bounded autonomy, polycentric governance, and value-aligned metrics, which require a lifecycle framework capable of handling high variability while preserving systemic integrity.

This section outlines the seven lifecycle states of a UME, the transition conditions between them, and the governance, support, and accountability protocols that apply at each stage. It also covers failure-state handling (degradation and collapse) as part of the lifecycle.

Lifecycle Phases of a UME

Phase	Description	Key Governance Trigger	Supporting Structures
1. Genesis	A proposed UME is formed around a mission, team, and purpose.	HAO or local SEP endorsement & alignment check	Charter template, DEA clause, Seed investment
2. Incubation	Early operational phase; feasibility, coherence, and commitment tested.	DEA v1 activation + resource allocation	Mentorship pod, embedded facilitator, VAM
3. Validation	Initial product/service-market fit is tested; internal metrics reviewed.	Internal milestone review + SEP potential	Peer audit, SEP readiness checklist
4. Maturation	UME becomes operationally stable; begins contributing to the HAO.	Revenue threshold + governance participation	Equity agreements, reduced HAO % take
5. Integration	UME joins SEPs or cross-UME initiatives; becomes a network contributor.	SEP contract or inter-UME agreement	Dynamic contribution models, inter-op registry
6. Evolution	UME reconfigures due to success, shifting market, or team evolution.	Proposal to modify DEA scope or governance model	Dynamic restructuring protocol
7. Dissolution	UME winds down, exits, or merges with another entity.	Voluntary, conditional, or collapse-triggered	Exit protocol, asset & equity redistribution

Detailed State Descriptions

1. Genesis

- **Inputs:** An individual or small group initiates a proposal, submitted to the HAO or a SEP (a joint venture between teams) formation unit.
- **Outputs:** UME Charter (mission, boundaries, values), draft DEA (a versioned operating agreement replacing fixed bylaws) section, preliminary needs assessment.

- **Dependencies:** A provide.io-like entity may offer tooling, financial modeling, or initial scaffolding.

2. Incubation

- **Purpose:** Establish working dynamics, define governance rhythms, test viability.
- **Conditions:** Access to seed resources; weekly governance pulse check; facilitator assigned.
- **Support Mechanisms:** Templates, onboarding rituals, embedded sociotechnical coach, ECC and VAM (ongoing checks that actions match stated principles) hooks enabled.

3. Validation

- **Purpose:** Test whether the UME can produce aligned value and sustain basic operations.
- **Triggers:** Metrics such as delivery consistency, internal accountability, peer trust thresholds.
- **Risks:** Cultural misalignment, burnout, role ambiguity.

4. Maturation

- **Purpose:** The UME begins full participation in the ICN (the reference cooperative business network) ecosystem.
- **Triggers:** Net positive value flow for 2+ cycles, governance participation, adherence to transparency thresholds.
- **Shifts:** Contribution to HAO revenue share drops (e.g., from 40% → 25%).

5. Integration

- **Purpose:** Cross-enterprise work with other UMEs and/or external entities.
- **Mechanisms:** Strategic Enterprise Partnership (SEP) formation, shared IP, co-investment logic.
- **Governance Note:** Requires a Multi-UME Joint Operations Agreement (MUJOA).

6. Evolution

- **Purpose:** The UME changes shape, for example by spinning out a sub-UME, pivoting its function, or reforming leadership.
- **Triggers:** Role fragmentation, market signal shifts, cultural drift, or scaling thresholds.
- **Tooling:** Dynamic restructuring templates, VAM-triggered adaptation prompts, member re-contracting tools.

7. Dissolution

- **Paths:**
 - *Voluntary Exit:* Mission fulfilled or members choose closure.
 - *Reintegration:* Merger into another UME.
 - *Conditional Dissolution:* Triggered by unmet governance or value thresholds.
 - *Collapse:* Emergency dissolution due to conflict, financial insolvency, or protocol violation.
- **Procedures:** Structured offboarding, equity redistribution via recovery protocol, open record-keeping.

Transition Triggers

Each lifecycle transition is governed by three types of triggers:

1. **Objective Metrics**
e.g., revenue flow, deliverable velocity, participation index, cultural health scores
2. **Governance Thresholds**
e.g., approval by internal quorum, SEP invitation, consent-based progression
3. **System-Level Conditions**
e.g., strategic redundancy needed, cross-network demand, resource reallocation

Failure Modes and Degradation Pathways

A UME may exhibit symptoms of degradation before collapse:

- Decline in value alignment (VAM threshold breach)
- Internal role atrophy or authority vacuum

- Sentiment analysis indicating dysfunction
- Repeated governance bypasses or quorum failures

Interventions may include:

- Temporary suspension of operations
- Injection of facilitation or governance support
- Triggering a network-level **Enterprise Recovery Review (ERR)**

If degradation proceeds to collapse, the UME’s:

- Knowledge assets are archived
- Trust and equity stakes are evaluated for restitution
- Members are invited to join other units or spin out new ones

This decentralized, coherent approach to failure handling is intended to support antifragility and cultural continuity during systemic stress.

Lifecycle Monitoring Tools

- **Value Alignment Monitoring (VAM)**
Tracks internal actions vs. declared principles
- **Enterprise Culture Cultivation (ECC)**
Measures cohesion, communication rhythms, and psychological safety
- **Lifecycle State Index (LSI)**
Computed indicator summarizing readiness for transition
- **Participant Sentiment Tracking**
Anonymized inputs on internal health

Conclusion

The lifecycle framework of the UME functions as a system architecture rather than a project-management model. By encoding structured emergence, permissioned evolution, defined exit paths, and structured failure-handling, the HAO framework aims to preserve network coherence without constraining autonomy. Lifecycle functions as a mechanism for learning, adaptation, and regeneration, in addition to survival.

§6.2 Strategic Enterprise Partnership (SEP) Lifecycle and Governance

Introduction

Strategic Enterprise Partnerships (SEPs) — joint ventures between teams — are temporary or long-lived collaborative ventures formed between two or more UMEs (United Micro Enterprises) — small, self-managing venture teams of up to ~15 people. They operate within the Humanized Autonomous Organization (HAO) — the network’s coordinating framework. Unlike UMEs, which are persistent, semi-autonomous value-producing nodes, SEPs are cooperative structures created to pursue a defined objective, address a network opportunity, or build shared infrastructure.

SEPs do not override UME autonomy; they provide a lightweight, accountable coordination layer. Their lifecycle is contractual, scoped, and versioned, often bound by shared outcomes, deliverables, and reinvestment logic. SEPs may span multiple UMEs, include HAO facilitation, and, when interfacing with external markets, may form the legal basis for external-facing entities (e.g., public market interfaces).

This section outlines the six-state lifecycle model of a SEP, its formation and dissolution mechanics, and the governance, resourcing, and value-flow protocols that apply throughout its existence.

SEP Lifecycle States

Phase	Description	Key Governance Trigger	Representative Outputs
1. Proposal	A need, opportunity, or shared goal is identified between two or more UMEs.	Joint intent expression + SEP charter draft	SEP Charter v0, preliminary budget, UME alignment
2. Formation	SEP is instantiated through consent-based ratification by participating UMEs.	SEP Charter v1 + Dynamic Contribution Agreement	Governance schedule, initial funding, resourcing map
3. Operation	SEP executes its chartered mission through coordinated activity.	Workstream launch + milestone tracking	Deliverables, shared infrastructure, service provisioning
4. Evolution	SEP adjusts scope, membership, or terms based on internal or external factors.	SEP Charter re-versioned with new alignment terms	Addenda to governance, resource reallocation

Phase	Description	Key Governance Trigger	Representative Outputs
5. Completion	SEP concludes its purpose or completes its deliverables.	Exit trigger met or project scope fulfilled	Handover artifacts, asset division, sunset report
6. Transition	SEP transitions assets, personnel, or IP into the broader HAO.	SEP wind-down protocol or handoff executed	Legacy protocol activation, archival, equity finalization

1. Proposal Phase

- **Initiators:** Two or more UMEs; may also include HAO catalysis or external opportunity signals.
- **Deliverables:**
 - Preliminary SEP Charter (intent, boundaries, contribution logic)
 - Initial risk/benefit matrix
 - Stakeholder mapping
- **Tooling:** SEP Charter templates, Contribution Planning Canvas, VAM (ongoing checks that actions match stated principles) pre-check

2. Formation Phase

- **Mechanisms:**
 - Consent-based ratification from participating UMEs
 - Optional third-party facilitator (HAO or provide.io-like entity)
- **Outputs:**
 - **SEP Charter v1**, a versioned governance contract
 - **Dynamic Contribution Agreement (DCA)**: maps input types to output entitlements
 - **Mutual Accountability Schedule**: outlines rituals, governance cadence, exit options

Note: Formation includes creation of a SEP Ledger Address (on the DLI), allowing for transparent resource tracking and interoperability with HAO-wide infrastructure.

3. Operation Phase

- **Activities:**
 - Joint execution of workstreams
 - Use of collaborative intelligence tools (CIN integration)
 - Asynchronous and synchronous governance rhythms
- **Governance Features:**
 - Lightweight decision-making (e.g., consent, rotating facilitator)
 - Ongoing value alignment checks (VAM active)
 - SEP participants may maintain dual roles in their home UMEs

Example Outputs:

- Open-source toolkits
- Shared service platforms
- Infrastructure used by multiple UMEs
- New market channels

4. Evolution Phase

- **Trigger Conditions:**
 - Market change
 - UME withdrawal or addition
 - Governance inefficacy or scope creep
- **Processes:**
 - SEP Charter re-versioned (v2, v3, ...)

- Adjustment of contribution agreements
- Addition/removal of stakeholders
- **Tooling:**
 - SEP Health Check
 - Collaborative renegotiation session
 - SEP Equity Map update

5. Completion Phase

- **Triggers:**
 - Objective fulfilled (e.g., infrastructure delivered)
 - Strategic value exhausted
 - Deliverables sunset per original charter
- **Outputs:**
 - SEP Completion Report (summary, value generated, lessons)
 - Retrospective with VAM metrics
 - Asset and surplus distribution logic
- **Financial Note:**
 - Final SEP revenue sharing is executed based on DCA or result-based equity formula

6. Transition Phase

- **Scenarios:**
 - Reusable infrastructure absorbed by HAO
 - Key contributors spin off a new UME
 - Public-facing SEP transitions to Public Market Interface (PMI) — a buffer company between the network and outside investors
- **Transition Tools:**
 - Asset custody transfer protocol
 - SEP Legacy Licensing (LL)
 - Member re-assignment or retirement logic
- **Governance Conclusion:**
 - Archive SEP ledger on DLI
 - Deactivate SEP governance rhythm
 - Optional: convert SEP into long-lived shared institution (e.g., HAO-scale service node)

SEP Governance Characteristics

Attribute	Model
Decision Protocol	Consent-based with escalation path
Charter Structure	Versioned, modular, identity-linked
Conflict Resolution	Escalation to HAO mediator or network-level arbitration pool
Participation Eligibility	Must be an active UME, SEP, or authorized HAO unit
Equity & Ownership	Defined per-contribution; SEP-specific, not persistent outside charter

Failure Modes and Emergency Protocols

SEPs may degrade or fail due to:

- Value misalignment
- Withdrawal of key participants
- Unresolvable coordination deadlock

Recovery/Exit Paths:

- Soft dissolution with asset handover

- UME-led SEP salvage and relaunch
- Emergency SEP Arbitration Review (SAR) from HAO

Lifecycle Monitoring Tools

- **SEP Ledger Snapshot:** Real-time resource and deliverable tracking
- **SEP Trust Health Index:** Aggregated participant sentiment and reliability score
- **SEP Contribution Graph:** Visualization of labor and material inputs
- **Charter Drift Detector:** Alerts when behavior deviates from chartered scope

Conclusion

Strategic Enterprise Partnerships function as crosslinking structures within the HAO ecosystem, connecting otherwise independent UMEs, enabling larger-scale efforts, and supporting emergent forms of inter-entity coordination. Their bounded, modular, versioned structure is intended to support flexibility, collaboration across entities, and shared value creation without requiring centralized control or loss of UME autonomy.

§6.3 Enterprise Collapse and Containment Protocols

Introduction

In decentralized socio-technical systems like the Humanized Autonomous Organization (HAO) — the network’s coordinating framework — the collapse of a United Micro Enterprise (UME) — a small, self-managing venture team ($\leq \sim 15$ people) — or a Strategic Enterprise Partnership (SEP) — a joint venture between teams — is treated as an expected, design-integrated event rather than an aberration. Collapse is part of the network’s adaptive lifecycle: a failure mode that also serves as an occasion to preserve integrity, recycle value, and support systemic resilience.

This section formalizes the triggers, stages, containment processes, and post-collapse actions that apply when a UME or SEP can no longer fulfill its governance commitments, deliver on its chartered intent, or maintain alignment with the network’s stated values. These protocols cover:

- Structured offboarding of members
- Containment of contagion or disruption
- Preservation of valuable artifacts (code, IP, relationships, capital)
- Continuous network coherence and resilience

Collapse vs. Dissolution: Definitions

Term	Description
Dissolution	A planned or voluntary wind-down, usually post-deliverable or mission-fulfilled
Collapse	A non-voluntary or condition-triggered failure state due to internal breakdown, financial insolvency, conflict, or value misalignment
Enterprise Degradation	A pre-collapse state characterized by drift from core principles, reduced functionality, or social dysfunction

Collapse Triggers

Collapse is activated when one or more of the following are detected:

1. Governance Failure

- Inability to reach quorum
- Repeated bypassing of decision protocols
- Consent deadlock or role abandonment

2. Value Misalignment

- Breach of the Dynamic Enterprise Agreement (DEA) — a versioned operating agreement replacing fixed bylaws
- Negative threshold breach in **Value Alignment Monitoring (VAM)** — ongoing checks that actions match stated principles

3. Operational Breakdown

- Unrecoverable financial insolvency
- Extended service delivery failure
- Persistent member attrition without recovery

4. Conflict Escalation

- Internal disputes escalated to HAO arbitration and unresolved
- Psychological safety breakdown (per ECC reports)

5. External Threat

- Legal, security, or reputational crisis impacting the broader HAO

Collapse Lifecycle

Phase	Description	Key Action
Detection	Collapse risk identified via automated metrics, member reporting, or HAO watch functions	Signal classification and internal review
Triage	Fast assessment of severity, risk domain (social, financial, reputational), and recoverability	Initiate Enterprise Recovery Review (ERR)
Intervention	Attempt to halt collapse with remediation, facilitation, or temporary stewardship	Assign emergency facilitator or SEP link
Confirmation	Collapse is declared if remediation fails or thresholds are crossed	Official network-level collapse broadcast
Containment	Quarantine access, suspend permissions, freeze assets, de-escalate stakeholder impact	Enact Containment Protocol Tier (CPT-1/2/3)
Deconstruction	Safely unwind agreements, redistribute assets, archive work, and initiate cultural healing	Trigger Recovery Protocols (RP-1/2)
Reintegration	Survivable fragments, members, or assets rejoin the HAO via vetted pathways	SEP assimilation, equity redistribution

Containment Protocols (CPT)

Containment Tiers are activated based on severity:

Tier	Trigger Conditions	Actions
CPT-1 (Soft Fail)	Low-impact degradation or misalignment	Suspend permissions, initiate internal mediation
CPT-2 (Medium Fail)	Governance breakdown or delivery failure	Freeze assets, trigger HAO facilitation, notify dependent SEPs
CPT-3 (Hard Fail)	Crisis escalation, breach of trust, or reputational threat	Total lockout, legal review, trigger HAO crisis cell

Each tier mandates a corresponding **Containment Response Packet (CRP)** containing:

- Authorized roles
 - Timeline for resolution
 - Member support contacts
 - SEP/UME fallback protocols
-

Deconstruction Steps

1. Asset Inventory

- Snapshot DLI ledger entries
- Identify shared or licensed IP
- Freeze SEP contracts and dependencies

2. Equity and Compensation Resolution

- Execute **Enterprise Recovery Protocol (ERP)** clause from DEA
- Determine member entitlements post-collapse
- Redistribute surplus or debt with VAM-informed weighting

3. Artifact Preservation

- Archive documentation, codebases, social graphs
- Mark assets as reusable, deprecated, or restricted
- Reassign stewardship if reusable

4. Member Reintegration or Exit

- Optional guided re-integration into new or existing UMEs
 - Restorative justice process for affected parties
 - Conflict debrief or cultural healing sessions
-

Post-Collapse Signals and Learning

Collapsed UMEs and SEPs are treated as learning artifacts for the network. The HAO performs:

- **Retrospective Autopsy Report (RAR):** Timeline of events, signals, and decisions
- **Cultural Resonance Check:** Did the collapse shift network norms, sentiment, or values?
- **Policy Adaptation Review:** Were existing protocols insufficient? What must evolve?

Reports are indexed and accessible via the HAO’s **Network Intelligence Ledger**, unless sealed due to sensitivity.

System-Level Safeguards

The following systems mitigate systemic risk from collapse:

- **Distributed Trust Graph:** Prevents overdependence on any one UME
- **Dynamic Equity Allocation:** Limits unearned exit value
- **Redundant Pathways for Critical Services:** SEPs or HAO can absorb mission-critical roles
- **HAO Crisis Cells:** Standby pods trained for rapid response and triage
- **Enterprise Contingency Simulation (ECS):** Periodic scenario drills

Conclusion

In HAO systems, collapse is treated as an expression of adaptability, boundary testing, and principled exit rather than an anomaly. What distinguishes a resilient network from a brittle one is its preparedness to respond, contain, and learn from collapse. The Enterprise Collapse and Containment Protocols are intended to preserve the integrity of the whole even as individual parts undergo transformation or dissolution.

§6.4 Governance Protocol Evolution

Introduction

In Humanized Autonomous Organizations (HAOs) — the network’s coordinating framework — governance is versioned, modular, and context-aware rather than static, monolithic, or centralized. It is designed to evolve alongside the cultural, operational, and strategic development of each UME (United Micro Enterprise) — a small, self-managing venture team (≤ ~15 people) — each SEP (Strategic Enterprise Partnership) — a joint venture between teams — and the broader HAO network.

This section outlines the theory, structure, and mechanics of governance protocol evolution within the HAO model. It describes how the **Dynamic Enterprise Agreement (DEA)** — a versioned operating agreement replacing fixed bylaws — changes over time, how versioning and modularity are preserved, and how decision-making around protocol shifts is made resilient, inclusive, and auditable. Unlike static bylaws or traditional amendments, governance in the HAO is treated as a living contract layer: responsive to feedback, aligned with network-wide principles, and designed to resist capture or drift.

Core Concepts

Term	Definition
Dynamic Enterprise Agreement (DEA)	A versioned, modular governance artifact defining operational, ethical, and legal structures
Protocol Evolution	The deliberate alteration of governance logic in response to network insight, strategic need, or emergent risks
Constitutional Layer	The set of principles and values (e.g., ETHICAL framework) that govern how governance itself may evolve
Nested Governance	The model by which each HAO entity (UME, SEP, HAO-core) maintains semi-autonomous but interoperable governance layers

Triggers for Governance Evolution

Governance protocol changes may be triggered by:

1. **Internal Feedback Loops**
 - Governance health checks
 - Repeated quorum failures or low participation
 - Sentiment analysis from members
2. **Operational Stress**
 - Collapse events (see §6.3)

- Emergent complexity beyond current rules
- Coordination overhead spikes

3. Environmental Shifts

- Regulatory change
- Market evolution
- Cultural divergence in stakeholder expectations

4. Network-Wide Initiatives

- Harmonization across multiple UMEs
- Experimental governance methods needing formalization
- Upgrades to cross-UME coordination protocols

Versioning Structure of the DEA

Each UME or SEP maintains its own **DEA instance**, consisting of:

- **Core Modules** (shared across network):
 - ETHICAL framework
 - Decision-making rules (e.g., consent, delegation)
 - Conflict resolution and escalation
 - Collapse protocols
- **Contextual Modules** (local to entity):
 - Role definitions
 - Resource allocation models
 - Compensation structures
 - Legal entity wrapper (LLC, coop, etc.)
- **Experimental Modules** (optional):
 - Piloted governance mechanisms
 - Alternate voting/consent schemas
 - Embedded AI-assisted decision systems

Each DEA is versioned using a semantic format (e.g., `v3.2.1`), where:

- Major versions signal fundamental change to principles or structure
- Minor versions include rule or policy updates
- Patch versions indicate clarifications, corrections, or tuning

Version history is immutably stored in the **Distributed Ledger Infrastructure (DLI)**.

Protocol Change Mechanisms

1. Proposal Phase

- Any member can initiate a Governance Change Proposal (GCP)
- Proposals must include:
 - Clear rationale linked to ETHICAL principles
 - Draft of affected clause(s)
 - Stakeholder impact assessment
 - Optional simulations or case studies

2. Review Phase

- **Nested review** by affected governance bodies:

- Within UME: role council, contributor pods
- Across SEPs: multi-UME review board
- At HAO-level: Protocol Integration Circle (PIC)
- Tooling:
 - Governance simulation engines (e.g., behavior trees)
 - Stakeholder polling
 - AI-supported sentiment classification

3. Consent Phase

- Decision thresholds vary by scope:
 - UME-local: internal consent or supermajority
 - SEP-spanning: quorum of participant entities + alignment check
 - Network-wide: HAO Consent Protocol (3-phase: proposal, deliberation, ratification)

4. Integration Phase

- Once approved:
 - DEA instance is re-versioned
 - Reference implementation published
 - Any affected runtime systems are updated (e.g., smart contracts, APIs, automation hooks)
 - Changes are annotated, timestamped, and linked to prior versions.
-

Evolution Constraints

Governance evolution is constrained by:

- **Constitutional Guardrails**
 - ETHICAL principles may not be removed or contradicted
 - Consent structures may not be replaced with coercive models
 - **Cross-Compatibility Rules**
 - Interoperability with other DEA instances must be maintained
 - Shared protocols (e.g., SEPs, VAM (ongoing checks that actions match stated principles) hooks) must remain forward-compatible
 - **Time-Based Moratoriums**
 - Prevent excessive volatility by requiring cooldowns between major revisions
-

Meta-Governance: Evolving the Evolution Rules

The HAO allows for recursive governance: the rules for changing governance can themselves evolve. To help ensure stability:

- Meta-protocol changes (e.g., how consent works) require **network-wide deliberation**
 - Multi-stakeholder modeling tools are required before vote eligibility
 - Proposals must include:
 - Cross-domain risk assessment
 - Change impact on nested governance layers
 - Fail-safe rollback procedures
-

Failure Modes and Recovery

If governance protocol evolution goes awry:

- **Rollback Protocol:** Previous DEA version can be reinstated if:
 - Post-change metrics degrade

- New protocol triggers unintended collapse behavior
 - **Quarantine Mode:** Allows experimental governance within a test-only context
 - E.g., pilot a new voting model within a UME for 3 months
 - **Meta-Stall Detection:** Monitors for governance gridlock or abuse of the protocol change process (e.g., fork loops, veto spam)
-

Auditability and Transparency

All DEA changes are:

- Logged to the DLI with:
 - Proposer, timestamp, affected clauses, rationale, simulation data
- Accessible via version diff tools
- Included in **Governance Transparency Reports (GTRs)** for members and auditors

Optional: AI-generated summaries for member accessibility, and public dashboards for SEP/UME comparison.

Conclusion

In a system designed to support ongoing emergence, distributed agency, and value-aligned innovation, governance must itself be governable. The evolution of governance in HAOs is treated as a core function rather than a disruption. Through structured versioning, constraint layering, consent mechanisms, and audit trails, the HAO is designed to support governance that remains responsive, transparent, and adaptive rather than fixed.

§6.5 Intentional Ephemerality

Introduction

Source and provenance. This section rests on a single 12-line fragment of an AI-assisted interview transcript, `~/code/site-provide-coop/Intentional Ephemerality.txt`. Direct inspection separates the opening prompt fragment (line 4), which asks whether dissolving a successful Cellular Venture could produce shared learning, from three Gemini response bullets (lines 8, 10 and 12), which elaborate the mechanisms below. The response ends by asking about rotation breaks; it does not establish them as a rule. No independent source or adoption record verifies these proposals as network policy. The section preserves them as exploratory design options.

Not every dissolution in the Humanized Autonomous Organization (HAO) — the network’s coordinating framework — follows from failure. This section describes units designed, from their founding, to end: a United Micro Enterprise (UME) — a small, self-managing venture team ($\leq \sim 15$ people) — or a Strategic Enterprise Partnership (SEP) — a joint venture between teams — whose charter fixes a finite lifespan, or a defined end-condition, from the outset, independent of whether the unit goes on to succeed.

`docs/06-lifecycle/01-ume-lifecycle-and-transition-states.md` already lists **Dissolution** as one of seven UME lifecycle phases, reached through voluntary exit, reintegration, conditional dissolution, or collapse. Intentional ephemerality describes a narrower, earlier design choice: the lifespan is fixed at **Genesis**, written into the unit’s charter or Dynamic Enterprise Agreement (DEA) — a versioned operating agreement replacing fixed bylaws — rather than decided later, once a mission is judged complete. The model’s material on this design pattern predates the current UME/SEP vocabulary and describes it at the level of a general venture rather than specifying UME or SEP; it applies most directly to the SEP, whose existing Completion phase (`docs/06-lifecycle/02-sep-lifecycle-and-governance.md`) already carries a sunset report and asset-distribution logic, but nothing in the model restricts the design to SEPs alone.

Distinguishing Designed Finitude from Collapse

This section is not about failure, and should not be read as a variant of `docs/06-lifecycle/03-collapse-and-containment-protocols.md`. That section covers **Enterprise Collapse**: a non-voluntary, condition-triggered failure state, reached through governance failure, value misalignment, operational breakdown, conflict escalation, or external threat, and followed by containment tiers (CPT-1 through CPT-3) and deconstruction steps. None of those triggers apply to a unit governed by intentional ephemerality. Its wind-down is scheduled, or condition-triggered from the start, rather than a response to drift or breakdown. A unit’s wind-down falls under exactly one of the two protocols in this chapter, never both, and the containment machinery described in §6.3 has no role in an ephemerality-driven wind-down.

Designed-Finite Ventures

Intentional ephemerality departs from the assumption that persistence is itself evidence of success. A unit’s finite lifespan is set deliberately, as a check against internal ossification and as a way of keeping resources in circulation rather than allowing them to settle into fixed positions. A unit may be wound down at the height of its success: the design decouples the end of a unit’s operating life from its performance, so dissolution under this section is a scheduled event rather than a diagnosis of unit performance.

The stated rationale is functional: it is intended to let the network redeploy resources, or end a unit that has outlived its usefulness, without waiting for a decline to make the case. A fixed lifespan is treated as one more parameter set at Genesis, alongside a unit’s charter, DEA scope, and initial resourcing, rather than as an outcome negotiated only once the unit is already underway.

The Graceful Sunset Process

Winding down a designed-finite unit is described in three parts, given together as a bundle rather than as separately sequenced steps:

1. Distribution of profits to members.
2. Dissemination of accumulated knowledge.
3. Encouragement of members to found new ventures addressing different problems within the network.

The model does not specify the sequencing, timing, or governance sign-off for these three actions, or which role executes them; this section does not supply a more elaborate procedure than the source gives. Where additional mechanism is wanted, it draws on the network’s existing dissolution machinery — structured offboarding, equity redistribution, and open record-keeping, per `docs/06-lifecycle/01-ume-lifecycle-and-transition-states.md`, or the SEP Completion phase’s handover artifacts and sunset report, per `docs/06-lifecycle/02-sep-lifecycle-and-governance.md` — rather than from any procedure specific to intentional ephemerality itself.

Leadership Assessed on Wind-Down, Not Persistence

The model reframes the metric by which leadership is judged: leadership is evaluated not by how long a unit persists, but by how effectively it distributes resources — both tangible and knowledge-based — before dissolving to seed new work. Facilitating a successful, graceful sunset is treated as standing alongside the launch of a new initiative, rather than as a lesser achievement. A stated implication is that new models of recognition are needed to reward this kind of leadership; no specific recognition mechanism is given.

A related, undeveloped framing positions sunset-at-success as a counter to entrenched too-big-to-fail dynamics: keeping resources in ongoing circulation rather than letting them accumulate inside units that have stopped needing them.

Rotation Breaks: A Proposed Mechanism, Not an Adopted Rule

The model raises, as an open design question rather than a stated policy, whether leadership positions should carry mandatory rotation breaks: a leader steps out of their seat and into a mentorship role elsewhere in the network, explicitly to prevent power from settling permanently in one area of expertise. This is preserved here as a proposal under consideration, because that is how the source material itself frames it — as a question, not a rule.

Several details are unspecified. Which leadership positions the proposal would cover — every leadership role network-wide, or only unit-level leads — is not stated. The break’s duration is not stated. Whether it would be triggered on a schedule, on a tenure threshold, or by some other condition is not stated. This section names the gap rather than filling it.

This is a different mechanism from the temporary rotation described for Mission-Driven units in `docs/06-lifecycle/06-lifecycle-variants-by-unit-type.md`: there, Operational-unit members rotate temporarily into a Mission-Driven unit to relieve burnout — a movement into a support role. Here, a leader rotates out of a leadership seat into mentorship — an anti-entrenchment measure. Both use the word “rotation,” but the two describe unrelated movements, in opposite directions, for different reasons, and should not be conflated.

The closest existing material in the corpus is incidental rather than a precedent for this proposal: a worked financial example lists rotating leadership stipends and rotating sabbaticals as options for allocating surplus (`docs/appendices/cafe-examples.md`), not as a governance mechanism, and it does not connect to the anti-entrenchment rationale described here.

Related but Distinct: Intentional Cannibalization

Intentional ephemerality should also be kept separate from **Intentional Cannibalization** (`docs/appendices/terminology.md` §VI), the pattern by which the network itself proactively dissolves or restructures a UME it judges outdated, to seed better-aligned successors. The two differ in who initiates the wind-down and when the decision is made: ephemerality is a unit’s own founding design, fixed before the unit does any work; cannibalization is a network-level judgment made about an existing unit, at a time of the network’s choosing.

Conclusion

Intentional ephemerality adds a founding-stage design option to the UME and SEP lifecycles: a unit whose end is scheduled rather than earned or forced. It shares vocabulary with, but is functionally separate from, both Enterprise Collapse (unplanned failure) and Intentional Cannibalization (a network-initiated judgment about an existing unit). The model gives a three-part sketch of what a graceful sunset involves, and a reframed metric for judging the leadership that carries one out; it leaves the mechanics of sequencing, timing, and any anti-entrenchment rotation requirement unspecified.

§6.6 Lifecycle Variants by Unit Type

Introduction

Source and provenance. This section rests on `~/code/site-provide-coop/provide.coop/gemini-on-clm.mdx`, three successive AI-generated passes over the same framework, staged for this corpus and never checked against any other source. The staged extraction describes it as a drafting artifact rather than a ratified specification, and notes that its passes are not always mutually consistent with each other — including on how many lifecycle stages the framework carries. Per-unit-type assignments below were checked against that primary file directly, because an earlier staged summary of the same material misattributed one of them. The typology is recorded because the corpus's own lifecycle treatment is uniform across unit types and this is the only source that varies it; it is not adopted as specification.

The seven-phase lifecycle described in `docs/06-lifecycle/01-ume-lifecycle-and-transition-states.md` — Genesis, Incubation, Validation, Maturation, Integration, Evolution, Dissolution — applies uniformly across every United Micro Enterprise (UME) — a small, self-managing venture team ($\leq \sim 15$ people). This section adds a differentiated layer on top of that uniform model: three unit types, each carrying a distinct operating posture, that call for different handling at several of the same lifecycle stages. It does not replace or renumber the seven-phase model; it describes how that model plays out differently depending on which kind of UME is passing through it.

A note on terminology. The three-type framework and the seven-phase UME lifecycle come from different points in the model's development and use their own stage vocabulary — proposal, formation, onboarding, monitoring, growth, and transition or dissolution — rather than Genesis through Dissolution. The two sequences cover similar ground closely enough to read side by side, but no source states a formal one-to-one correspondence between them, and this section does not assert one. A stage discussed below should be read as an elaboration of the territory a like-named UME phase covers, not as a renaming of that phase.

The Three Unit Types

Unit type	Core definition	Operating posture
Operational	Units core to the network's everyday functions — finance, legal, onboarding, and similar standing work	Requires stability, predictability, and a risk-averse approach; described in the source material as the network's backbone
Mission-Driven	Units addressing a specific market need or social problem	Thrives on adaptability and measured innovation; carries the stated pressure of balancing ethical practice against market viability
Research and Exploratory	Units focused on open-ended questions, including potentially high-risk projects	Prioritizes a culture of learning and iteration; granted the widest latitude of the three types

Type-Differentiated Treatment Across the Lifecycle

Proposal

Feasibility review at the proposal stage takes the intended unit type into account. Proposals for Operational units emphasize long-term viability. Proposals for Research and Exploratory units include an explicit statement of the learning goal the unit pursues, alongside its stated potential benefits.

Formation

Standard formation templates carry type-specific addenda — temporary leadership roles for Research and Exploratory units, for example. Dynamic Enterprise Agreement (DEA) — a versioned operating agreement replacing fixed bylaws — clauses can include pre-planned exit provisions specific to type: an escape hatch for a Research and Exploratory unit that reaches an ethical gray area, and a sunset provision for a short-term Mission-Driven unit. This is the one point in the model where a sunset mechanism is tied to a specific unit type rather than offered network-wide; see `docs/06-lifecycle/05-intentional-ephemerality.md` for the network's general treatment of a unit designed from Genesis to have a finite lifespan.

Onboarding

Onboarding is type-specific in both content and connections. Operational units train on reliable process execution; Mission-Driven units train on conflict resolution; Research and Exploratory units train on treating failure as productive rather than as a discouraged outcome. Cross-unit introductions follow the same logic: Operational units are linked with innovative ventures in similar industries, and Research and Exploratory units are connected with community groups that stand to be affected by their work.

Monitoring and Support

The same signal reads differently by type: high variance in a unit's progress can be a positive indicator for a Research and Exploratory unit and a warning sign for an Operational unit. Support is tailored per type:

- **Operational** units receive heavy investment in mentorship, so that reliable processes pass on and the unit's dependence on any specific individual falls.
- **Mission-Driven** units receive **ethical stress-testing** alongside the standard viability checks applied to every unit, in recognition of the stated pressure to balance ethical practice against market performance. Support at this stage can include temporary rotation of Operational-unit members into the Mission-Driven unit, to relieve burnout. This is a distinct mechanism from the leadership rotation breaks described in `docs/06-lifecycle/05-intentional-ephemerality.md`: there, a leader rotates out of a seat and into mentorship elsewhere in the network; here, members from a different unit type rotate in temporarily to provide relief.

- **Research and Exploratory** units are granted the widest latitude of the three types, offset by a mandatory cross-unit reporting practice — **Show and Tell** — intended to prevent that latitude from becoming isolation and to allow early, network-wide detection of a research direction headed somewhere harmful.

Growth

Growth is also read differently by type. Operational-unit growth is judged on reliability over speed. A successful Mission-Driven unit is expected to spawn semi-autonomous sub-units to preserve agility as it grows. Findings from a Research and Exploratory unit can lead to an entirely new unit type being recognized within the network.

Transition or Dissolution

Reaching this stage remains, across all three types, a healthy possibility rather than a default indicator of failure. An Operational unit may be absorbed into the network's core function rather than dissolved outright. A Mission-Driven unit's dissolution is oriented toward maximizing the lessons learned for similar future ventures. Across all three types, the model applies one diagnostic question network-wide at this stage: whether a unit's lifespan was a natural conclusion or a signal of a systemic flaw, examined for patterns across unit types rather than case by case.

A Support Function Requiring Composite Expertise

Supporting units at this level of differentiation is described as requiring expertise spanning data analysis, organizational dynamics, and forward-looking judgment, rather than a single specialization. Success at this stage is also measured through signals — sentiment analysis of internal communication, anonymized well-being surveys — that sit alongside, rather than inside, the network's existing Value Alignment Monitoring (VAM) — ongoing checks that actions match stated principles — and Enterprise Culture Cultivation (ECC) tracking. VAM and ECC track alignment and cohesion; the stress- and burnout-specific signals described here are not currently named as part of either.

What This Section Does Not Cover

This typology is a layer added to the UME lifecycle model, not a substitute for it. It does not use, and should not be read as reviving, the earlier “Enterprise Degradation” or “Enterprise Collapse” vocabulary; failure-triggered dissolution, for any unit type, remains governed by `docs/06-lifecycle/03-collapse-and-containment-protocols.md`. Nor does the differentiation described here extend to Strategic Enterprise Partnerships (SEPs) — a joint venture between teams — whose own lifecycle is described separately in `docs/06-lifecycle/02-sep-lifecycle-and-governance.md`; the model's source material for this typology precedes the current SEP/UME distinction and speaks only in terms of the unit type, without extending it to partnerships between units.

Conclusion

Operational, Mission-Driven, and Research and Exploratory units pass through the same named lifecycle differently: different proposal emphases, different formation addenda, different onboarding content, different monitoring signals and support, different growth patterns, and different framings of a healthy transition. None of this changes the seven-phase model itself; it specifies how that model should be read differently depending on which of the three types a given UME is.

7.0 • Human Systems

7.0 Introduction

The Humanized Autonomous Organization (HAO) — the network’s coordinating framework — inverts a common assumption in organizational design: it treats technology and governance as instruments that serve members, rather than the reverse. Many organizations optimize primarily for efficiency, scalability, and shareholder return; HAOs are designed to prioritize member well-being, collective agency, and participatory governance. This section describes how HAOs aim to support human needs across emotional, cognitive, social, and existential dimensions through embedded design patterns, relational protocols, and socioemotional infrastructure.

The human-centered dimension of the HAO is treated as structural rather than incidental. Financial, operational, and technological interactions are treated as relational touchpoints: opportunities to build trust, reinforce shared meaning, and support mutual accountability. This differs from models that optimize primarily for transactional efficiency or behavioral engagement metrics.

Human-centered design (Norman, 2013), research on workplace engagement and psychological safety (Kahn, 1990), and trauma-informed organizational change (Bloom & Farragher, 2013) inform this section. Applying those bodies of work to HAOs is a corpus design synthesis; the architectural and procedural elements are intended to help HAOs remain resilient and inclusive as they grow.

Key components include:

- **Socioemotional Safety and Shared Language:** Psychologically safe environments where members can participate openly without fear of reprisal or exclusion.
- **Cultural Onboarding and Ritual Design:** Initiation practices that extend beyond compliance training to support alignment with organizational values and identity.
- **Conflict Engagement as Organizational Literacy:** Non-punitive, adaptive mechanisms for addressing tension, misalignment, and difference.
- **Regenerative Learning Systems:** Systems that transfer knowledge and cultivate adaptive intelligence, intergenerational insight, and communal memory.

Together, these practices reflect a design orientation toward organizations that develop people as well as economic output. The model treats members as agents embedded in interdependent systems, rather than as interchangeable units.

This section describes design patterns intended to support human dignity, empathy, and purpose within high-tech, highly autonomous environments.

References

- Bloom, S. L., & Farragher, B. (2013). *Restoring Sanctuary: A New Operating System for Trauma-Informed Systems of Care*. Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780199796366.001.0001>
- Kahn, W. A. (1990). Psychological conditions of personal engagement and disengagement at work. *Academy of Management Journal*, 33(4), 692–724.
<https://doi.org/10.5465/256287>
- Norman, D. A. (2013). *The Design of Everyday Things: Revised and Expanded Edition*. Basic Books.

§7.1 Socioemotional Safety and Shared Language

Human participation in an organizational system depends on logic and incentives, and on emotional safety, relational clarity, and perceived belonging. In HAOs (the network’s coordinating framework), where authority and participation are distributed, socioemotional safety supports coherence. Without it, governance and economic systems are vulnerable to unspoken tensions, interpersonal distrust, and cultural drift.

This section defines socioemotional safety in HAOs and describes the role of shared language protocols in supporting it at scale.

7.1.1 Definition of Socioemotional Safety

Socioemotional safety refers to the felt experience that one’s thoughts, emotions, identity, and contributions can be expressed and acknowledged without fear of ridicule, reprisal, or exclusion. It is a context that supports participatory risk-taking, which HAOs treat as a condition for innovation, accountability, and collective intelligence.

It differs from conflict avoidance or interpersonal ease. Socioemotional safety supports differentiated expression, principled disagreement, and the surfacing of tacit knowledge, provided these are expressed within culturally held norms.

Key conditions:

- Emotional expression is normalized, not pathologized
 - Vulnerability is met with attunement, not avoidance
 - Mistakes are treated as learning opportunities, not punishable offenses
 - Power imbalances are named and contextualized, not hidden
-

7.1.2 Structural Enablers of Socioemotional Safety in HAOs

Socioemotional safety is systematized through the following mechanisms:

- **Norm-Setting Rituals:** Every UME (a small, self-managing venture team, $\leq \sim 15$ people) and SEP (a joint venture between teams) engages in regular re-grounding of shared behavioral expectations through reflection protocols (e.g., working agreements, conflict norms, tension check-ins).
 - **Check-in/Check-out Practices:** Meetings include structured emotional presence rituals that calibrate tone and allow members to surface unseen dynamics early.
 - **Consent-Based Governance:** By privileging “no strong objection” over majority rule, decision-making processes encourage dissent and include edge perspectives.
 - **Role-Based Distinctions:** Clearly articulated roles depersonalize critiques during feedback cycles.
 - **Third-Party Emotional Facilitation:** Neutral facilitation is available for tension mediation, emotional labor load-balancing, or cultural drift detection.
-

7.1.3 Shared Language as Protocol

Language is the interface layer of a human system. Without a shared baseline, meaning-making can fragment across the network.

HAOs establish **shared language protocols**, structured vocabularies that:

- Codify values (e.g., empathy, integrity, autonomy)
- Draw on nonviolent-communication norms
- Normalize naming tension, power, and emotions directly
- Provide escalation and de-escalation rituals (“I need a pause”, “Let’s call a meta-conversation”)

Examples:

- **“Check myself”:** A phrase signaling self-reflection rather than defense
- **“Needs not met”:** A non-blaming way to raise unmet expectations
- **“From a place of care...”:** Preceding difficult feedback with intent

Shared language is taught during onboarding and reinforced during retrospectives, peer reviews, and governance meetings. Over time, it functions as organizational middleware: it standardizes communication across the network without enforcing uniformity.

7.1.4 Psychological Safety Across Scales

HAOs operate in nested layers: individuals within UMEs, UMEs within SEPs, and SEPs within the broader HAO. Each layer requires context-specific applications of socioemotional safety:

Level	Application
Individual	Clear role expectations, personal check-ins, access to mentorship
UME	Norms for internal disagreement, shared language rituals, retrospective facilitation
SEP	Cross-team cultural calibration, tension resolution protocols, nonviolent negotiation spaces
HAO	Meta-governance accountability processes, emotional debt tracking, distributed ombudsperson roles

7.1.5 Anti-Patterns and Failure Modes

Without deliberate cultivation, the absence of socioemotional safety can lead to:

- **Shadow Hierarchies:** Informal power dynamics replacing transparent roles
- **Avoidance Cultures:** Fear of raising tensions, which can suppress innovation
- **Burnout Cycles:** Emotional labor disproportionately falling on unrecognized members
- **Cultural Fragmentation:** Misalignment between subgroups without mechanisms for integration

HAOs treat these as structural indicators of system misalignment, not individual failings.

7.1.6 Conclusion

Socioemotional safety functions as an operational input in HAOs. By embedding shared emotional norms, codifying language practices, and distributing emotional labor across the network, HAOs support participation at scale without centralized control.

§7.2 Cultural Onboarding and Ritual Design

7.2.1 Introduction

Traditional onboarding processes focus on information transfer: handbooks, workflows, compliance training. In contrast, HAOs (the network’s coordinating framework) require onboarding that transmits culture, values, and shared mental models—closer to initiation than orientation. Cultural onboarding in HAOs is an iterative, relational, and participatory process that brings new members into coherence with the network’s ethical and operational rhythms.

Alongside onboarding, ritual design serves as the protocol layer for collective meaning-making. Rituals in HAOs are functional, repeatable practices that encode governance, emotional safety, accountability, and renewal.

Together, cultural onboarding and ritual design form a **social membrane**—a semi-permeable boundary that integrates people into the system without losing coherence or over-assimilating diversity.

7.2.2 Functions of Cultural Onboarding

In the HAO context, onboarding serves four functions:

1. **Alignment** – Deep understanding of HAO principles (e.g., subsidiarity, surplus reinvestment, consent governance)
2. **Contextualization** – Understanding how a UME (a small, self-managing venture team, $\leq \sim 15$ people) or SEP (a joint venture between teams) fits into the broader HAO ecosystem
3. **Trust Building** – Early-stage relational integration with team members, mentors, and governance processes
4. **Protocol Literacy** – Familiarity with shared language, norms, escalation paths, and socioemotional safety tools

Rather than an HR task, onboarding functions in HAOs as a rite of passage—a transition from outsider to trusted participant.

7.2.3 Onboarding as a Layered Process

HAOs use a multi-layer onboarding design, calibrated to match the network’s nested structure:

Layer	Onboarding Focus
HAO-Wide	ETHICAL + PARTS frameworks, governance orientation, protocol awareness
UME/SEP-Specific	Operating rhythm, local rituals, decision rights, role structure
Role-Specific	Expectations, accountabilities, feedback mechanisms, tools used
Relational Layer	Mentorship, buddy systems, onboarding cohorts, reflection sessions

Each layer has distinct timing (e.g., initial entry, 30-day, 90-day, 6-month cycles), and feedback from new members is used to assess cultural clarity.

7.2.4 Ritual Design in HAOs

Rituals are the operational encoding of culture: embodied, repeatable practices that formalize shared intentions and anchor distributed action in emotional coherence.

Examples include:

- **Weekly Pulse Meetings:** Structured time for collective sensing—“How is the team doing? What’s unsaid?”
- **Governance Retrocycles:** Regular cycles where team members reflect on governance friction points, evolving protocols accordingly
- **Conflict Rituals:** Pre-agreed scripts for naming tensions, pausing interactions, and inviting facilitation
- **Gratitude Rounds:** Scheduled expression of appreciation to surface invisible contributions and reduce emotional debt
- **Onboarding Ceremonies:** Initiatory moments where new members are welcomed, express intent, and are witnessed in their commitment

The design of these rituals adheres to the following principles:

- **Opt-in by default**, with transparent purpose
 - **Flexible in form**, fixed in function
 - **Anchored in time**, recurring with intention
 - **Designed with feedback loops**, to evolve over time
-

7.2.5 Anti-Patterns and Cultural Drift Risks

When onboarding is minimal or rituals are performative:

- New members experience cultural dissonance—uncertainty about what’s safe or expected
- Important values remain symbolic, not enacted
- Norms are passed through informal cliques, creating shadow onboarding and misaligned expectations
- Decision-making slows due to ambiguous protocol literacy
- Turnover can increase as members fail to internalize belonging

HAOs address these risks by treating onboarding and ritual design as evolving infrastructure, not static processes.

7.2.6 Cultural Feedback Mechanisms

Each onboarding cohort feeds back into the cultural system through:

- **Structured debriefs** at 30-, 60-, and 90-day intervals
- **Mentor reflections** logged in the governance system
- **Anonymized pattern sensing**, analyzing onboarding outcomes across time
- **Ritual retrospectives**, evaluating whether key ceremonies still serve their intended function

This reflexivity is intended to keep rituals and onboarding adaptive, not rigid or outdated.

7.2.7 Conclusion

In HAOs, culture is ritualized, internalized, and enacted. Cultural onboarding is designed for new members to integrate into the network’s practices, beyond formal membership. Ritual design gives the system rhythm, emotional coherence, and the ability to evolve without fragmenting.

By investing in these layers of human infrastructure, HAOs aim to translate values from principle to practice, extending each new member both responsibility and belonging within the network.

§7.3 Conflict Engagement as Organizational Literacy

7.3.1 Introduction

In conventional organizations, conflict is often treated as a threat—to productivity, authority, or cultural harmony. In Humanized Autonomous Organizations (HAOs) — the network’s coordinating framework — conflict is treated as a signal of misalignment, unmet needs, systemic incoherence, or emergent transformation. Rather than suppressing or escalating conflict, HAOs cultivate conflict engagement as organizational literacy—a shared capacity distributed across the network.

This section describes how HAOs process conflict systemically, using it to inform protocol changes and rebuild trust. The goal is alignment achieved through principled disagreement, not unanimous agreement.

7.3.2 Framing Conflict as Information

HAOs view conflict as:

- **Diagnostic:** reveals tension between values, incentives, or roles
- **Mirror:** reflects deeper structural or emotional patterns
- **Catalyst:** enables system learning and relational depth

This reframing repositions conflict from a breakdown to a feedback channel—like stress testing in resilient infrastructure.

To engage conflict well, HAOs aim to reduce:

- **Interpersonalization:** “You are the problem”
- **Avoidance loops:** Deferral until rupture
- **Punitive escalation:** Skipping reflection in favor of retribution

Instead, HAOs use structured spaces to surface, process, and route tension into protocol changes.

7.3.3 Structural Supports for Conflict Engagement

To support conflict literacy at scale, HAOs implement a multi-tiered resolution architecture:

Tier	Conflict Type	Protocol
0	Internal tension (intra-personal)	Personal reflection tools, journaling, peer mirroring
1	Dyadic misalignment	Nonviolent communication scripts, peer facilitation
2	Role-level or project misfit	Role negotiation spaces, facilitated retrospectives
3	Patterned team dysfunction	Circle processes, anonymous pattern mapping, ritual pauses
4	Systemic or values-level misalignment	Governance proposals, value audits, advisory review councils

HAOs offer multiple entry points—decentralized but traceable—each paired with ritualized de-escalation practices.

7.3.4 Embedded Conflict Protocols

Tools used in day-to-day conflict engagement include:

- **Tension Tracking Boards:** Members can anonymously or publicly post tensions, tagged by area (relational, process, governance, equity).
- **Nonviolent Communication (NVC) Scaffold:** “When X happened, I felt Y because I need Z. Would you be willing to...?”
- **Role Clarification Sessions:** Invoked when feedback consistently maps to unclear accountability.
- **“Calling a Meta”:** Any member can pause a discussion to name an unspoken tension or emotional drift.
- **Emotional Facilitator Rotation:** Roles responsible for naming affective undercurrents in recurring meetings.
- **Grievance → Transformation Pipeline:** Formal complaints are routed into a reflection + redesign loop rather than adversarial hearings.

These practices are intended to lower the barrier to naming conflict, reduce shame and ambiguity, and model responsible ownership of discomfort.

7.3.5 Culture of Conflict Fluency

Conflict literacy is cultivated through:

- **Onboarding rituals** that include emotional vocabulary and tension-mapping exercises
- **Regular conflict simulation labs** (e.g. roleplays, “difficulty drills”, resolution games)
- **Feedback-in-action rounds**, where live responses to disagreement are examined post-meeting
- **Conflict closure rituals**, honoring the end of a conflict cycle and releasing residual affect

HAOs develop ritualized language to normalize and de-shame conflict:

- “Let’s metabolize this tension together”
 - “I think there’s a design signal underneath this friction”
 - “I’m willing to stay in the discomfort if you are”
 - “This feels like a misalignment of interpretation, not intent”
-

7.3.6 Misuses and Failure Modes

Even with strong design, HAOs can fall into failure patterns:

- **Therapy culture creep:** Confusing personal healing with collective processing mandates
- **Feedback weaponization:** Using feedback for status assertion or unprocessed grievance
- **Conflict exceptionalism:** Avoiding recurring low-level tension in favor of dramatic rupture cycles
- **Process fatigue:** Over-formalizing responses to minor issues, leading to disengagement

Mitigating these risks requires balance: enough structure to protect safety, enough flexibility to allow organic repair.

7.3.7 Conflict as Strategic Capacity

In HAOs, conflict serves several functions:

- It surfaces dissent relevant to decisions
- It can reveal boundary conditions relevant to evolving protocols
- When processed, it is intended to support mutual trust
- It surfaces misalignments that might otherwise go unnoticed

A network’s capacity to engage conflict directly, without collapse or avoidance, is treated as one indicator of organizational maturity.

7.3.8 Conflict Practitioner Roles

A separate body of process material names four practitioner roles that recur across conflict-engagement work, each attached to a different point in the process rather than to a single case type:

Role	Function
Conflict Coach	Works one-on-one with a single party, preparing them for a difficult conversation before it takes place.
Facilitator	Structures direct dialogue between the parties to a conflict, without deciding the outcome.
Mediator	Runs formal mediation for complex, escalated disputes, as an impartial third party with no decision authority.
System Designer	Designs and revises the conflict-resolution processes themselves, rather than working individual cases.

The material these roles are drawn from does not itself reference the tier structure in §7.3.3, and none of the four roles is assigned to a tier by that source.

One cross-reference is added here as an editorial observation, not as a sourced claim: the Facilitator’s stated function — structuring direct dialogue between two parties without deciding the outcome — matches, in wording, this chapter’s own Tier 1 protocol of peer facilitation for dyadic misalignment (§7.3.3). The System Designer role is explicitly cross-case by its own description, operating on the resolution processes themselves rather than on individual disputes, so it is not tied to any single tier for that same source-stated reason. For the Conflict Coach and Mediator roles, no tier correspondence is stated in the source or inferred here.

7.3.9 Conclusion

Conflict engagement in HAOs is a designed function intended to support system health, individual accountability, and collective learning. By treating conflict as literacy rather than liability, HAOs aim to make disagreement a route to alignment, not a threat to it.

§7.4 Regenerative Learning Systems

7.4.1 Introduction

In conventional organizations, learning is typically episodic, externalized, or transactional—limited to upskilling, compliance training, or managerial development. In Humanized Autonomous Organizations (HAOs) — the network’s coordinating framework — learning is **regenerative**: continuous, embedded, and participatory. It is designed to build capability across people, roles, and time, in addition to accumulating knowledge.

A regenerative learning system is designed to produce:

- Individual growth as well as performance
- Organizational memory as well as metrics
- Collective adaptability as well as alignment

Learning in HAOs functions as a network-wide capability, integrated into governance, conflict resolution, onboarding, compensation, and dissolution.

7.4.2 Definitions and Learning Modalities

Regenerative learning refers to learning that:

- Sustains and enhances individual and system health
- Produces more learning (meta-learning)
- Is distributed, not centralized
- Is triggered by work, not separate from it

Learning modalities in HAOs include:

Modality	Description
Experiential	Reflection and adaptation based on doing (e.g., retrospectives, role reviews)
Peer-based	Horizontal learning from cohort or circle-based knowledge exchange
Mentorship	Relational transmission of values, norms, and tacit knowledge
Emergent	Learning that surfaces from breakdowns, conflicts, or edge cases
Reflexive	Learning about learning—identifying how the system learns, forgets, and remembers

7.4.3 Embedded Learning Structures

HAOs embed learning through recursive scaffolds:

- **After Action Reviews (AARs)**: Standardized debriefs following initiatives, sprints, or conflicts that feed into operational adjustments
- **Live Protocol Tuning**: Policies, norms, or tools are revisited in real time during use, as well as during scheduled reviews
- **Pattern Libraries**: Codified learnings made accessible in lightweight digital systems (e.g., “what worked in past onboarding cycles”)

- **Living Role Cards:** Role definitions evolve based on reflections, peer feedback, and real-world friction
 - **Cross-UME/SEP Knowledge Syncs:** Horizontal learning events or asynchronous knowledge transfers across UMEs (self-managing venture teams) and SEPs (joint ventures between teams)
 - **Learning Pods:** Small rotating learning groups focused on emergent topics (e.g., facilitation skills, conflict fluency, new toolchains)
-

7.4.4 Intergenerational Knowledge Transmission

Most systems are optimized for speed and output; few are designed for continuity and memory.

HAOs treat organizational knowledge as intergenerational capital:

- Departing members leave behind learnings via structured exit interviews and memory deposits
- Long-serving members serve as culture anchors as well as senior contributors
- Rituals of storytelling, myth-making, and foundational failures are preserved and surfaced at key moments
- Governance evolution is accompanied by change logs and value-based rationales, in addition to version control

This is intended to make evolution distinct from repetition, with learning treated as longitudinal, not episodic.

7.4.5 Metrics of Regenerative Learning

Rather than measuring completion rates or certifications, HAOs track:

- **Protocol velocity:** Frequency and success rate of policy or tool adaptations
- **Feedback loop fidelity:** How often signals from the edge inform core decision-making
- **Distributed facilitation index:** How many members take active roles in retros, tuning, and learning events
- **Organizational amnesia risk:** Presence of single points of failure in knowledge domains
- **Relearning rate:** How often past problems resurface without acknowledgment of history

These indicators track whether the organization is improving its judgment as well as accumulating information.

7.4.6 Learning Culture and Incentives

Learning is treated as a shared responsibility and is reinforced through:

- **Time allocation:** Dedicated capacity for personal and collective learning is built into workload planning
- **Incentives:** Contributions to learning (facilitation, documentation, mentorship) are recognized in compensation frameworks
- **Cultural signals:** Learning from failure is normalized through visible leadership modeling
- **Rituals:** Learning reviews are ritualized (e.g., quarterly “Harvest Weeks,” story circles, failure feasts)

This positions learning as infrastructure for resilience and evolution, rather than an optional extra.

7.4.7 Conclusion

HAOs treat learning as a network function, not an individual pursuit. Regenerative learning is intended to help the system improve over time: to observe itself, revise itself, and recover from disruption. By embedding learning at every level of operation—ritual, tool, contract, role—HAOs aim to avoid combining speed without depth, or complexity without coherence.

In this model, the network aims to adapt under uncertainty without abandoning its stated principles or purpose.

§7.5 Synthesis — Human Systems as Infrastructure

7.5.1 Reframing the Infrastructure Stack

In traditional organizational design, “infrastructure” typically refers to physical systems (buildings, networks, logistics) or digital systems (APIs, databases, cloud services). Human systems—onboarding, emotional safety, conflict resolution, and learning—are often treated as soft skills or organizational culture: abstract and optional.

In a Humanized Autonomous Organization (HAO) — the network’s coordinating framework — this framing is inverted: human systems are treated as infrastructure, on the same basis as software or physical systems, designed with attention to version control and redundancy, similar to software platforms or

distributed ledgers. The human stack—language, ritual, emotional norms, sensemaking protocols—is treated as part of core operations, supporting scale without hierarchy and change without collapse.

This section summarizes the components described earlier and describes how HAOs approach human-system design as an intentional, adaptive architecture.

7.5.2 The Integrated Human Infrastructure Stack

Across the HAO, human systems form an interoperable infrastructure layer, composed of:

Layer	Core Function	Mechanism
Emotional Safety Layer	Prevents fragmentation	Socioemotional protocols, check-ins, role containment
Cultural Encoding Layer	Maintains coherence	Ritual design, onboarding scaffolds, symbolic containers
Conflict Metabolism Layer	Enables feedback and adaptation	De-escalation protocols, multi-tier resolution systems
Learning Nervous System	Generates intelligence	Reflexive practices, knowledge transmission rituals
Trust Signaling Layer	Regulates access and risk	Progressive verification, role boundaries, relational capital
Ethical Governance Layer	Aligns action with principles	Dynamic agreements, representation systems, veto points

Each layer is composable and recursive—able to function independently but optimized in relationship. They are designed to support both operational throughput and recovery of system health under stress.

7.5.3 Principles of Durable Human Infrastructure

To support long-term viability, HAOs design human systems using principles adapted from civil, software, and ecological infrastructure design:

- **Idempotence:** Interactions should be repeatable without degradation
- **Observability:** Tensions, patterns, and breakdowns must be visible and legible
- **Degeneracy:** Multiple different systems can perform overlapping functions (e.g., learning via mentorship or retrospectives)
- **Intentional Inefficiency:** Some processes (like consensus building) are slow by design, to support alignment
- **Versionability:** Norms, rituals, and agreements must be updatable with traceable rationale
- **Intergenerational Capacity:** Systems must preserve knowledge beyond current participants

These design principles treat human experience as data to be interpreted and incorporated, not friction to be eliminated.

7.5.4 Toward a General Theory of Human-System Engineering

This chapter describes early patterns for a discipline referred to here as **human-system engineering**. In this paradigm:

- Culture is designed rather than emergent
- Conflict is metabolized rather than managed
- Learning is endogenous, not off-platform
- Governance is internalized as infrastructure, not treated only as external compliance

This shift requires that people who build HAOs work with emotional data, symbolic systems, collective intelligence, and ethical coherence as design material.

7.5.5 Conclusion

In this model, the scalability and adaptability of HAOs depend on both technology and engineered human systems—rituals, protocols, norms, and roles that encode values and operationalize trust. They are the infrastructure on which the model’s autonomous, cooperative networks are built.

By treating emotional health, relational trust, cultural coherence, and adaptive learning as infrastructure, the HAO model treats these systems as a distinguishing design feature relative to conventional organizations.

8.0 • Legal

The viability of a Humanized Autonomous Organization (HAO) — the network’s coordinating framework — depends on more than internal logic, cultural alignment, or technological sophistication. Deployment requires legal structures and regulatory strategies that encode autonomy, enforce contracts, and manage liability while maintaining decentralization, ethical integrity, and human primacy.

This section presents a modular legal architecture for HAOs and their constituent units (e.g., United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people; Strategic Enterprise Partnerships (SEPs) — joint ventures between teams; Member Trust Unions (MTUs) — the network’s credit-union-like financial institutions). It addresses the tension between organizational innovation and legal conformity by proposing jurisdictionally adaptable templates, flexible governance bindings, and mechanisms to translate post-conventional governance into enforceable legal terms.

HAOs differ from traditional firms and cooperatives in structure and in their rejection of centralized authority as a prerequisite for legitimacy. This makes legal integration complex, requiring a multi-layered framework that encodes nested sovereignty, trust-driven finance, and evolving agreements without undermining local autonomy.

Key focuses of this section include:

- Modular entity design (e.g., LLCs, co-ops, B-corps, DAO-LLCs) suited for UMEs and the HAO itself
- Legal strategies for nested governance and distributed liability
- Commons-based intellectual property models
- Public interface translation mechanisms (e.g., for interfacing with traditional investment structures)
- Regulatory sandboxing, jurisdiction selection, and legal failover mechanisms

We assume a mixed-jurisdictional deployment model in which UMEs may exist across legal boundaries, and in which HAO networks function legally as both a federated system and a single economic actor. Beyond compliance, this section considers legal templates that could inform how law treats novel organizational forms.

This section provides templates for governing documents, strategies for shared liability buffers, and patterns for contractual interoperability with conventional organizations and capital systems.

§8.1 Modular Legal Forms and Multi-Level Structures

Enabling legal recognition of autonomous economic actors within federated governance systems

Overview

The Humanized Autonomous Organization (HAO) — the network’s coordinating framework — requires a legal substrate flexible enough to express interdependencies across distributed units while preserving autonomy, accountability, and lawful operation across jurisdictions. Unlike traditional firms with a single legal personality, HAOs are composed of semi-autonomous agents: United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people; Strategic Enterprise Partnerships (SEPs) — joint ventures between teams; and coordinating entities such as the HAO Core. Each may require distinct legal representation based on its function, size, and regional legal environment.

This subsection proposes a multi-level, composable legal architecture. It supports:

- Individual UMEs registering as single-purpose legal vehicles (LLCs, cooperatives, or community interest companies),
- The HAO itself as either a nonprofit holding company, cooperative trust, or networked foundation,
- SEPs as joint ventures with temporary, flexible legal scaffolding, and
- Dynamic agreements that bind these parts through governance protocols in addition to legal contracts.

8.1.1 UME Legal Forms (Edge Units)

UMEs are the smallest legally distinct economic actors in the HAO. Depending on mission, market, and jurisdiction, UMEs may take on various legal structures:

Form	Use Case	Strengths	Constraints
LLC (Limited Liability Company)	Service-producing UME, flexible operations	California guidance describes a flexible worker-cooperative vehicle that can include members outside the US	Cooperative governance and mission must be specified in an operating agreement, subject to local law
Worker Cooperative	Labor-owned ventures	Member governance and potential patronage-based surplus allocations	Entity and tax treatment depend on jurisdiction and structure; no general single-tax rule
Nonstock Nonprofit Corp (e.g., WI Ch. 181)	Mission-driven service units	Member governance is possible; federal tax exemption requires separate qualification	No routine patronage distributions; member payments and dissolution assets face statutory and, where applicable, tax-exemption limits
B-Corporation	Mission-aligned product ventures	Investor-friendly, reputation gain	Requires profit motive, ongoing certification
CIC (Community Interest Company) [UK]	Community-aligned UMEs	Legal lock on community purpose	Limited in scope, UK-specific
DAO-LLC (e.g., Wyoming, UT)	On-chain governance UME	Smart contract governance, experimental legitimacy	Legal novelty, uncertain judicial precedent

The California LLC and Wisconsin nonstock examples are drawn from research-library documents `icn-0037` and `icn-0040` in the references appendix. They are illustrations, not current legal templates. [California's LLC operating-agreement statute](#), [Wisconsin Chapter 181](#), and [IRS cooperative-return instructions](#) show why governance, distributions and tax treatment need separate jurisdiction-specific review.

Sexton's historical guide discusses both equal and investment-weighted votes, as well as admission and withdrawal provisions, within an LLC operating agreement (`icn-0037` , PDF p. 5). For the HAO model, choosing the LLC form alone would not establish equal member control; the governing agreement would need to specify voting and exit rules. The guide's older California filing and tax details are not used as current requirements.

A given HAO implementation — for example, ICN (the reference cooperative business network) or MTU (the network's credit-union-like financial institution) — may provide templates and infrastructure for registering UMEs under these forms with embedded governance bindings (e.g., "constitution as smart contract," or multi-sig key custodianship for compliance).

8.1.2 SEP Legal Scaffolding

SEPs are multi-UME collaborations that often need a lightweight, dissolvable legal form accommodating:

- Joint resource pooling
- Shared risk and reward
- Time-boxed or goal-bound operation

Proposed legal structures:

- **Contractual Joint Venture (JV LLC)**: Default pattern in the U.S. for co-managed entities
- **Multi-stakeholder Coop**: For SEPs with ongoing community engagement
- **DAO wrapper** with multi-UME governance (e.g., multisig between UME reps)
- **Purpose Trusts**: For SEPs with asset-holding or IP management goals

These may be instantiated using template generators or legal automation tooling developed by integrator organizations such as [provide.io](#).

8.1.3 The HAO Core Entity

The central coordinating layer of the HAO — responsible for core governance, infrastructure, and network-wide reinvestment — requires strong legal liability protection and the ability to interface with both edge units and external systems.

Recommended legal formations:

- **Cooperative Trust or Stewardship Foundation** (e.g., Swiss Verein, Dutch Stichting)
- **Nonprofit LLC or Hybrid Nonprofit + For-Profit Bifurcation**
- **Federated DAO LLC**, when working within DAO-recognizing jurisdictions
- **Holding Coop with Multi-Class Memberships** (for UMEs, contributors, investors)

This core entity should be:

- Structured to limit its own profit-taking from UMEs
- Able to enforce Dynamic Enterprise Agreements (DEAs) — a versioned operating agreement replacing fixed bylaws
- Able to own or license shared infrastructure (e.g., DLI, CIN)
- Capable of managing cross-border compliance with minimal administrative overhead

8.1.4 Inter-Entity Legal Agreements

To maintain cohesion across diverse legal units, DEAs bind entities via:

- Version-controlled constitutional documents
- Interlocking arbitration and mediation clauses
- Equity, revenue, and liability sharing protocols
- Protocol-first contract logic (optionally mirrored in legalese)

Each UME and SEP is legally independent but functionally interdependent via these agreements. In practice, this mirrors how federated systems like Mondragon maintain group-wide cohesion without collapsing into centralization.

8.1.5 Legal Portability and Templates

To streamline global HAO expansion:

- Core legal templates should be maintained in versioned repositories

- Local legal counsel should fork and adapt templates based on regional constraints
- All governance models should include “legal fallback paths” in the event of jurisdictional incompatibility (e.g., migration from LLC to Coop)

Tools like OpenLaw, LexDAO, and API-based entity registration services may be integrated to automate portions of legal lifecycle management.

§8.2 Jurisdictional Flexibility and Nesting

Enabling legal coherence across distributed units operating under heterogeneous regulatory conditions

Overview

Given the translocal and polycentric nature of Humanized Autonomous Organizations (HAOs) — the network’s coordinating framework — jurisdictional flexibility is a foundational requirement. Unlike traditional multinational corporations, which typically adopt a parent-subsidiary model under a centralized legal jurisdiction, HAOs consist of semi-autonomous legal entities — UMEs (small, self-managing venture teams of up to ~15 people), SEPs (joint ventures between teams), and MTUs (the network’s credit-union-like financial institutions) — operating across diverse geographies, each subject to local law, regulatory environments, and cultural norms.

This section outlines the nested legal architecture that allows a HAO to scale globally while retaining local legitimacy, legal safety, and structural coherence. It also introduces patterns for jurisdictional nesting, soft law harmonization, and decentralized compliance strategies.

8.2.1 Nested Legal Entities in a Federated Topology

HAOs typically use a nested legal entity model, where:

- UMEs operate as independent legal shells (LLC, coop, DAO, etc.) embedded within a larger federated network
- SEPs are often legally anchored in one of the participating UMEs’ jurisdictions or registered in neutral innovation zones
- The HAO Core acts as a meta-entity, holding governance protocols, IP, shared infrastructure, and reinvestment vehicles

These entities are bound through versioned Dynamic Enterprise Agreements (DEAs) — a versioned operating agreement replacing fixed bylaws — and interoperable legal templates, allowing for cross-entity recognition of decision rights, ownership stakes, and liability agreements.

Example:

A UME in Brazil operates under Brazilian cooperative law, a UME in Germany under GmbH law, and both form a SEP registered in Estonia as a DAO-compliant LLC. The HAO Core operates as a Swiss Verein or Dutch Stichting, managing shared contracts and equity instruments.

8.2.2 Jurisdiction Selection Criteria

When choosing jurisdictions for entity formation, HAO components should optimize for:

Factor	Considerations
Cooperative/DAO-Friendly Law	Jurisdictions like Vermont, Colorado (US), Estonia, Switzerland, and Wyoming are favorable
Taxation and Withholding	Minimize double-taxation and maximize pass-through options
Recognition of Digital Governance	DAO-LLC structures or smart-contract governance enforceability
IP and Data Protection	GDPR compliance, IP commons compatibility
Dispute Resolution Options	Availability of alternative dispute resolution (ADR) frameworks and arbitration clauses
Political Stability	Risk-adjusted analysis of legal continuity

A strategic deployment may involve registering the HAO Core in a trust-stable international jurisdiction, while allowing UMEs to choose home-ground legal shells suited to their operational realities.

8.2.3 Soft Law and Harmonization Layers

To ensure network-wide governance cohesion, HAO implementations rely on soft law protocols such as:

- Binding governance frameworks encoded in versioned DEAs
- Cross-jurisdictional recognition clauses (similar to multinational treaty recognition)
- Pre-registered arbitration protocols to handle inter-entity disputes
- Portable governance codebases, agreed upon via shared hash commitments

This creates a meta-legal layer: a de facto constitution respected across jurisdictions, even where full harmonization is not possible in hard law.

8.2.4 Legal Failover and Migration Patterns

When legal incompatibility, regulatory restrictions, or collapse conditions occur, entities must have:

- Failover protocols to migrate operations to another jurisdiction
- Legal forks of their agreements that remain valid even during transition
- Liquidation fallback provisions (e.g., preserve community-held equity, port governance state)
- Optionally, sovereign legal infrastructure via on-chain contracts or embedded logic (e.g., Ethereum-based arbitration)

Example:

If a UME's host country criminalizes DAO structures, its members may fork to a pre-approved cooperative template, while the HAO Core updates governance links and capital accounts accordingly.

8.2.5 Composability and Legal Standardization

To enable automated compliance and interoperability:

- Templates and schemas should be written using legal markup standards (e.g., OpenLaw, Accord Project's Cicero)
- Governance agreements should be digitally notarized and cross-signed by parent network nodes
- Entity metadata should follow a shared ontology, allowing federated discovery and audit

Pattern:

```
UME::legal_form: "LLC"  
UME::jurisdiction: "US-NY"  
UME::DEA_version: "v1.3.0"  
UME::sep_links: [SEP_ID#23982, SEP_ID#5543]
```

This creates a machine-readable registry of legal bindings and operational structures, enabling automated lifecycle tracking, status alerts, and audit readiness.

Conclusion

Jurisdictional flexibility is a design requirement for federated autonomy as well as a compliance obligation. Legal nesting, soft law harmonization, failover protocols, and composable templates are intended to support global viability while honoring local law and cultural specificity.

§8.3 Commons-Based Intellectual Property and Licensing Models

Designing commons-based IP frameworks for federated innovation systems

Overview

In conventional firms, intellectual property (IP) is typically owned by a centralized entity and monetized through exclusive rights, artificial scarcity, and legal enforcement. A Humanized Autonomous Organization (HAO) — the network's coordinating framework — takes a different approach to IP ownership and use.

Instead, HAOs treat IP as part of a shared knowledge commons: a resource that is co-created, co-maintained, and co-governed. This approach is intended to preserve contributors' rights while keeping innovation accessible, forkable, and value-aligned.

This section defines the HAO's approach to IP as a governable commons, offering patterns for attribution, governance, protection, and monetization of collective knowledge assets while limiting enclosure by any single platform or party.

8.3.1 IP in a Federated System

HAOs generate diverse forms of IP:

- Software systems (e.g., ledger tech, governance tooling, AI integrations)
- Business models and agreements (e.g., Dynamic Enterprise Agreements, templates for SEPs — joint ventures between teams)
- Cultural protocols and methodologies
- Design patterns, UI/UX systems, or data schemas
- Collectively maintained datasets

Unlike centralized firms, these IP forms emerge across multiple UMEs (small, self-managing venture teams of up to ~15 people) and SEPs, often co-produced by contributors in different jurisdictions. This requires multi-party attribution, modular licensing, and dynamic governance.

8.3.2 Default Licensing Strategy: Forkable but Governed

The baseline HAO IP policy follows a “forkable but values-aligned” model:

- Default license: **Cooperative Non-Compete License (CNCL)** or **Peer Production License (PPL)** or a modified **AGPL** with a value-alignment clause
- Explicit permission for commercial use only within aligned ecosystems
- Required attribution and share-alike terms
- Optional economic reciprocity (e.g., revenue sharing, coop dues)

These licenses:

- Prevent third parties from privatizing community-created IP outside the license terms
 - Incentivize aligned SEPs and UMEs to build on existing assets
 - Help contributors retain downstream leverage in hybrid public/private interfaces
-

8.3.3 IP Custodianship and Legal Stewardship

Rather than “owning” IP, the HAO Core or a designated **IP Steward Entity** serves as custodian of shared IP. This entity:

- Maintains license registries and versioning
- Approves exceptions or derivative licensing (e.g., commercial API access)
- Resolves attribution conflicts or derivative disputes
- Handles defensive IP strategies (e.g., preemptive patenting, CC0 reservation)
- Enables cross-jurisdictional recognition of common rights

The **IP Custodian** may be:

- A nonprofit legal entity (e.g., Stichting)
- A DAO-based rights registry with embedded governance logic
- A multi-UME representative board

This stewarding model mirrors existing commons-based organizations like the Creative Commons Foundation or GNOME Foundation, adapted for federated ownership and governance.

8.3.4 Attribution, Provenance, and Versioning

To maintain equitable recognition and governance:

- All HAO contributions should be version-controlled, attributed, and traceable via cryptographic signatures or contributor metadata
- Governance protocols should require collective attribution agreements on SEP deliverables or shared tooling
- Forks or derivatives must maintain transitive provenance metadata for governance and reward alignment

Example: A SEP co-develops a logistics framework (licensed under PPL). Any derivative UME-specific implementation must retain attribution to the SEP and use the same license unless granted an exemption via the IP Custodian.

8.3.5 Economic Use and Reciprocity Agreements

To allow commercialization while limiting one-sided value capture, HAOs may use:

- **Value-Aligned Usage Agreements (VAUA)**: terms defining how external organizations may license or integrate IP
- **Reciprocity tiers**: economic return scales with the licensee’s size, revenue, or usage
- **Dual licensing**: internal (member) and external (commercial) use models
- **SEP-specific licensing**: revenue generated from SEP-created assets may be governed under co-maintained smart contracts or time-bound use grants

This preserves IP as a financially generative, community-aligned asset.

8.3.6 Integration with External Licensing Models

To support legal interoperability:

- All HAO-compatible licenses should be OSD-compliant or Creative Commons-compatible where applicable
 - Contributor agreements should include pre-specified fallback licenses (e.g., MIT, AGPL) in the case of custodian collapse
 - HAO agreements should be cross-compatible with:
 - OpenChain (ISO/IEC 5230)
 - Open Source Hardware licenses
 - European Union Public Licenses
 - Data Commons frameworks (e.g., ODbL, CDLA)
-

8.3.7 Commons Degradation and Protection Mechanisms

To prevent enclosure, HAOs include:

- **Commons Watch Protocols:** alerting for misuse or enclosure
- Revocation rights in case of license violations
- Community enforcement clauses and arbitration paths
- Defensive IP registrations to prevent patent trolling

In edge cases, the HAO Core may trigger fork + exclusion mechanisms if a UME repeatedly violates commons terms.

Conclusion

By treating IP as a shared, governed commons, HAOs aim to support innovation while limiting enclosure and unilateral appropriation of shared assets. Licensing, stewardship, and economic reciprocity mechanisms form a core layer of this approach, intended to let federated networks scale without sacrificing coherence or values.

§8.4 Public Market Interaction Framework

Interfacing federated networks with traditional capital markets without compromising autonomy, equity, or ethics

Overview

One of the challenges facing Humanized Autonomous Organizations (HAOs) — the network’s coordinating framework — is engaging with external capital markets without compromising internal values. Traditional investment structures prioritize control, liquidity, and short-term returns, which are often incompatible with the HAO’s trickle-out economics and distributed-authority governance.

This section defines a buffered, selectively porous interface for HAOs to engage with public markets while preserving network integrity, local sovereignty, and member equity. The framework introduces **Public Market Interfaces (PMIs)** — a buffer between the network and outside investors. A PMI is a hybridized legal vehicle (e.g., a Contribulo-like structure) that translates between cooperative economics and traditional investor expectations.

8.4.1 Design Goals for External Engagement

Public market interactions must be designed to:

- Protect the **Micro Enterprise Ecosystem (MEE)** — the network’s protected internal economy — from pressure to divert value to outside investors
 - Maintain HAO majority ownership and governance rights
 - Translate between long-term value creation and short-term investor timelines
 - Allow external capital inflow without compromising internal equity dynamics
 - Provide financial return channels that align with ethical, social, and environmental goals
-

8.4.2 Public Market Interfaces (PMIs)

A PMI is a purpose-bound legal entity (LLC, PBC, or coop-hybrid) created to:

- Interface with capital markets on behalf of the HAO
- Serve as an equity wrapper or licensing body for one or more SEPs (joint ventures between teams)
- Distribute returns to investors under modified terms
- Maintain internal agreements with the HAO Core regarding value flow, licensing, and governance limits

Example:

In the model, Contribulo is a PMI that licenses software built by UMEs (small, self-managing venture teams of up to ~15 people) to external B2B clients. Its ownership allocation and any outside-investor share remain pending an owner decision; the example contemplates capped investor returns. Profits return to the ICN network — the reference cooperative business network — via licensing fees and revenue-sharing agreements.

8.4.3 Ownership and Governance Guardrails

To preserve sovereignty and alignment:

- PMIs must include HAO-aligned supermajority ownership (typically >66%)
- External investor shares may be:
 - Non-voting
 - Time-limited
 - Return-capped or subject to *earn-outs*
- PMIs should include veto protection clauses and mission-lock articles in governing documents
- IP licensed to PMIs must remain under revocable, values-aligned terms

8.4.4 Investor Return Structures

HAOs may use several financial instruments to satisfy investor expectations without compromising structure:

- **Royalty models:** Investors earn a fixed percentage of revenue from a PMI or SEP
- **Convertible revenue-share notes:** Repay principal plus a capped return over time
- **Tokenized dividends:** Blockchain-based payout mechanisms with embedded ethics guardrails
- **Return floors and ceilings:** Pre-defined return corridors with automatic payout redistribution to the HAO if exceeded

These mechanisms align incentives with the HAO's trickle-out economic model while mitigating risk for early investors.

8.4.5 SEP Integration and Revenue Routing

PMIs may represent one or more SEPs. To preserve traceability and fairness:

- Each SEP must define its own revenue allocation structure, including:
 - Contributions by participating UMEs
 - Investment from HAO Core or external partners
 - Resource utilization (e.g., HAO infrastructure)
- PMI agreements must respect the SEP charter, with enforceable fallback clauses if terms are violated

Revenue flows in three directions:

1. To participating UMEs (profit share, dividends)
2. To the HAO Core (network fees, reinvestment)
3. To investors (capped returns or royalties)

Diagram suggestion (not included here):

PMI as a center node routing value between public investors, SEPs, and internal HAO structures via bounded flows.

8.4.6 Market Signaling and Brand Strategy

HAOs and their PMIs must be transparent in their positioning:

- Clearly communicate the network’s value-flow model to potential investors
- Use impact metrics alongside financial KPIs to drive reporting
- Publish Public Engagement Charters that define acceptable forms of participation
- Avoid language that implies arbitrage opportunities or outsized speculative returns

PMIs may also serve to protect the HAO’s brand, limiting external engagements that could erode its mission, culture, or public identity.

8.4.7 Regulatory and Securities Compliance

Where PMIs involve share offerings or pooled investment vehicles, HAOs must:

- Ensure SEC/FINMA/EU-compliant disclosures if operating in regulated markets
- Use limited public offerings (e.g., Reg CF, Reg D, Reg A+) to test engagement models
- Consult securities attorneys to structure non-transferable share classes or ethics-based securities instruments

Legal structures such as **Trustless Common Shares** or **Participatory Income Shares** may be piloted where jurisdictionally permitted.

8.4.8 Failure Modes and Ethical Firewalls

In the event of mission drift, pressure to divert value to outside investors, or unsolicited takeover attempts:

- The HAO Core must retain the right to terminate PMI licensing agreements
- PMIs may be dissolved or forked under protective governance clauses
- UME- and SEP-level stakeholders must retain claim rights to derivative assets
- A recovery pathway must be designed into all PMI agreements, similar to social cooperatives’ mission-lock clauses

Pattern:

If a PMI attempts to IPO without HAO consent, all upstream licenses are revoked, revenue claims invalidated, and cooperative reinvestment withheld.

Conclusion

HAOs are unlikely to remain entirely isolated from public markets if they aim to scale. Engagement occurs on the HAO’s terms: buffered, principled, values-aligned, and reversible. PMIs function as a boundary object between two economic worlds, intended to preserve network integrity while providing access to capital, distribution, and visibility.

9.0 • Adversarial Analysis

Adversarial analysis examines how the mechanisms of a Humanized Autonomous Organization (HAO) — the network’s coordinating framework — can be turned against the purposes those mechanisms are specified to serve. The subject is the model’s own machinery: consensus procedures, mutual-support systems, value-alignment vocabulary, algorithmic decision support, and the distribution of authority across partly independent units. Each is a surface an actor can work against, and in several of the patterns cataloged here the property that makes a mechanism function is the same property that makes it available for misuse. The chapter is diagnostic. It names attack patterns, states the conditions under which each becomes available, describes what each would look like from outside, and records which of them the source material leaves unanswered.

Scope

The chapter covers five areas, which are also the five categories the threat catalog uses. §9.1.1 defines them and fixes their names; the names below are those names.

- **Mission and values** — patterns in which the network’s stated purpose is eroded by increments, hardened into criteria that cannot be tested, or invoked selectively against members (§9.1).
- **Governance capture** — patterns in which decision-making influence accumulates outside the channels specified to carry it, and patterns in which the specified procedure is used against its own purpose (§9.2).
- **Framework capture** — the case in which a unit acts against the network’s stated purpose and justifies the action in the network’s own vocabulary, so that the justification remains internally coherent on the network’s own terms (§9.3).
- **External and technical vectors** — co-option applied from outside the network, together with attacks on the shared technical systems and shared vocabulary the network depends on (§9.4).
- **Polycentric structural failure** — failure modes that follow from distributing authority across many partly independent centers, which the model inherits by adopting that structure (§9.5).

Detection and resistance mechanisms are cataloged in §9.6, and problems that no source resolves in §9.7. Affirmative governance design — how decision bodies are composed, how proposals pass, how disputes between units are resolved — belongs to chapter 03 and is cross-referenced here rather than restated.

Relationship to Existing Chapters

Four existing sections specify machinery that the threat model in this chapter applies to.

- `docs/06-lifecycle/03-collapse-and-containment-protocols.md` specifies what happens once a United Micro Enterprise (UME) — a small, self-managing venture team of up to ~15 people — or a Strategic Enterprise Partnership (SEP) — a joint venture between teams — is already failing: five classes of collapse trigger, tiered containment protocols, deconstruction steps, a retrospective autopsy report, and system-level safeguards such as the distributed trust graph and standby crisis cells. It is the response machinery. This chapter describes the attack patterns that lead to those triggers being reached.
- `docs/10-evaluation/02-participation-quality-and-alignment-audits.md` specifies participation auditing: consent decision audits, role health assessments, equity-engagement correlation mapping, and network-level engagement distribution. A number of the observable indicators listed in this chapter are quantities that auditing machinery already collects.
- `docs/10-evaluation/04-ai-augmented-governance-monitoring.md` specifies the Collaborative Intelligence Network (CIN) as AI-augmented rather than autonomous governance support, across three layers: human-in-the-loop deliberation aids, anomaly detection and scenario simulation (centralization drift, exit clustering, governance fatigue signals, SEP imbalance), and alignment monitoring against the network’s stated values. It is the standing detection layer. This chapter supplies the concrete patterns that layer would be configured to look for.
- `docs/15-advanced/06-systems-simulation-and-stress-testing.md` specifies agent-based modeling of unit, member, and network-level agents, failure-mode and antifragility scenario modeling — trust collapse in a regional Member Trust Union (MTU) — the network’s credit-union-like financial institution — governance fragmentation from versioning conflicts in the operating agreement, model misalignment producing decision drift, and reputation hoarding or sybil attacks in trust networks — and game-theoretic modeling of incentives and strategic behavior. It is the quantitative layer. This chapter catalogs qualitatively the behaviors that layer would simulate.

In each case the existing section specifies a mechanism, and this chapter supplies the threat model against which that mechanism is measured. None of the four is restated here.

Method and Its Limits

The material derives from red-team exercises conducted against the model’s own documentation. Three families of source contribute. Two adversarial interview passes were run against the Integrated Cooperative Network (ICN) — the reference cooperative business network — under prompts asking how the network could be ruined and how it could be destroyed (`gemini_icn-interview-1_p12r01_ruin-the-icn.md` and `gemini_icn-interview-1_p13r01_destroy-the-icn.md` , referred to below as p12r01 and p13r01). A single worked document, `~/code/papr/icn/security-bad-actors-and-mmm.md` , develops one scenario in which a unit exploits a legal gap and defends the act in the network’s own terms. Four model variants in `~/code/papr/polycentric-governance/` respond to a deliberately solution-free prompt asking for gaps in polycentric governance, and the same four respond to a follow-on prompt asking for solutions to those gaps; §9.5 draws on the first set and §9.6.20 on the second. All three families were consolidated into a single staged extraction before drafting, and the claims in this chapter trace to that extraction rather than to independent review of the archives.

Five limits follow from that provenance, and the first is the chapter’s principal limitation.

1. **No operating experience underlies any entry.** No unit described in this chapter has been observed. Every pattern is analytical: a statement of what the documented mechanisms would permit, derived from documentation rather than from deployment. No claim is made about how often any pattern occurs, how likely it is, or how damaging it would be in practice.
2. **The catalog is not closed.** It is the union of what these particular exercises produced under these particular prompts. Absence of a pattern from §9.1 is evidence about the source material, not about the model.
3. **Threat names are the sources' own.** They are retained so that later work can be traced back to the material that produced them, including where a name carries rhetorical framing. Where two sources use one name for unrelated things, or separate names for closely related things, both are recorded and neither is reconciled (§9.1.2, §9.1.6).
4. **One source contains constructed dialogue.** `security-bad-actors-and-mmm.md` presents a unit's self-justification as an authored illustration, and labels it as such. It is a construct, not testimony. No quotation attributed to a person appears anywhere in this chapter, and §9.3 treats the scenario as the analytical construct it is.
5. **The polycentric material is general governance literature.** The four variants converge on findings that public-administration research had already established about polycentric systems. §9.5 uses that material for its taxonomy of failure modes and marks it as inherited rather than as a finding about this model specifically.

§9.1 Threat Taxonomy

This section fixes the vocabulary the rest of chapter 09 uses. It assigns every attack pattern recorded in the staged extraction to one of five categories, catalogs the patterns in a single table, and then treats in detail the three that belong to the mission-and-values category. Threat names are the sources' own. They are retained even where a source's phrasing is rhetorical, because renaming them would break traceability to the material that a Humanized Autonomous Organization (HAO) — the network's coordinating framework — would consult if a pattern were suspected. Later sections refer to these patterns by the names given here and do not redefine them.

9.1.1 Categories

Five categories organize the chapter. The first four hold named attack patterns and are distinguished by what the pattern targets rather than by who conducts it, since several patterns are available to members and to outside parties alike. The fifth holds no named attacks.

- **Mission and values** — the network's stated purpose is the target. The pattern erodes it, hardens it into criteria that cannot be tested, or turns a single principle against the members who invoke the others. Detailed in §9.1.3 to §9.1.5.
- **Governance capture** — decision-making influence is the target. The category covers two routes to it, and §9.2 treats both. In the first, formal procedure continues to run while the influence that determines outcomes accumulates outside the channels specified to carry it, as in Hidden Oligarchy and Founder Capture. In the second, the specified procedure is itself used against its own purpose, so that following the rule produces the outcome the rule was adopted to prevent, as in Weaponized Consensus and Inaction as a Weapon. Detailed in §9.2.
- **Framework capture** — the network's legitimating vocabulary is the target. An action contrary to the network's stated purpose is defended in the network's own terms, and the defense holds together on those terms. Detailed in §9.3.
- **External and technical vectors** — the network's boundary and its shared infrastructure are the targets, including its legal environment, its public distinctiveness, its algorithmic decision support, and its shared vocabulary. Patterns in this category are grouped by target, so an attack on shared language conducted by an internal faction appears here rather than under governance capture. Detailed in §9.4.
- **Polycentric structural failure** — no attacker is required. These are consequences of distributing authority across many partly independent centers, and the model inherits them by adopting that structure. They are enumerated as eight failure modes in §9.5 rather than as entries in the catalog below, because the source material presents them as properties of a governance form rather than as patterns an actor executes.

9.1.2 Threat Catalog

Threat	Category	Mechanism	Section
Mission Creep	Mission and values	incremental concessions made under growth pressure accumulate until the network's founding commitments no longer constrain practice	§9.1.3
Purity Trap	Mission and values	escalating demands for ideological adherence convert stated values into admission criteria that cannot be tested or contested	§9.1.4
Mission Zealot	Mission and values	one stated principle is invoked selectively to suppress critique or to penalize performance	§9.1.5, with its governance expression in §9.2
Hidden Oligarchy	Governance capture	units with aligned interests coordinate informally outside official channels until they determine outcomes without holding formal authority	§9.2
Cult of Personality	Governance capture	a member converts consensus norms into personal loyalty and recasts critique of a decision as an attack on a person	§9.2
Weaponized Consensus	Governance capture	inclusivity norms are used to extend deliberation on low-stakes matters indefinitely, exhausting participants and blocking decisions	§9.2
Tyranny of Small Kindness	Governance capture	mutual-support capacity is absorbed by repeat claimants whose demands carry a social cost for anyone who questions them	§9.2
Founder Capture	Governance capture	a founder or long-committed member converts accumulated informal authority into control of the leadership pipeline and reframes dissent as disloyalty	§9.2

Threat	Category	Mechanism	Section
Weaponized Burnout	Governance capture	task allocation is skewed to overload the most values-aligned members until they disengage, clearing positions for others	\$9.2
Hijacking the Support Network	Governance capture	manufactured crises create dependence on informal helpers, who accrue influence outside the official support structure	\$9.2
Inaction as a Weapon	Governance capture	superficial compliance combined with minimal substantive contribution clogs decision channels and spreads distrust	\$9.2
Framework-turned-against-itself (loophole exploitation)	Framework capture	a unit exploits a legal gap and justifies the act through the network's own commitments to members, to legality, and to efficient resource use	\$9.3
Death by a Thousand Papercuts	Framework capture	many small, individually defensible acts of rule-bending accumulate below the threshold at which any single one can be proven	\$9.3
Regulatory Capture	External and technical vectors	incumbents lobby to reshape the legal environment against the network's structures, under framings drawn from unrelated protections	\$9.4
Greenwashing Through Imitation	External and technical vectors	conventional firms adopt the network's language and surface practices without its structures, blurring the public distinction between the two	\$9.4
Talent & Idea Poaching	External and technical vectors	competitors recruit governance-skilled members and copy the network's practices without the commitments those practices depend on	\$9.4
Divide and Conquer	External and technical vectors	outside parties circulate unverified accounts of hoarding or withheld knowledge to degrade trust between units	\$9.4
Trojan Horse	External and technical vectors	tools presented as efficiency improvements carry a bias toward financial return, and recommendations drift away from stated values without an audit that would reveal it	\$9.4
Poisoning the Well	External and technical vectors	training data is altered so that model outputs favor short-term return or discount conflict-resolution risk factors	\$9.4
Exploiting Blind Faith in Tech	External and technical vectors	misleading inputs are supplied to resource-allocation and partner-selection systems whose outputs are trusted without independent checking	\$9.4
Manufacturing Dissent	External and technical vectors	network-health monitoring is manipulated to display patterns of unfair distribution that the underlying records do not support	\$9.4
Lexical Drift	External and technical vectors	a faction shifts the connotation of shared terms over time until the vocabulary no longer constrains the behavior it was adopted to constrain	\$9.4

Alternative names in the sources. Several entries carry a second name in the material. Mission Creep is also called “The Death of a Thousand Compromises”. Hidden Oligarchy appears as “Silent Partner” in the earlier response of its source and as “Hidden Oligarchy” in the later two. Founder Capture appears as “Idealist Turned Cynic” and, in a more developed variant, as “Disillusioned Savior”. Lexical Drift is named “Capturing the Narrative” in its source, which identifies the target as the Mycorium, the sources’ term for the network’s evolving shared language and history. Inaction as a Weapon is paired in its source with “Social Engineering Attacks”. The primary name in the table is the one used throughout chapter 09.

Entries that overlap. Purity Trap and Mission Zealot describe closely related dynamics, and Hidden Oligarchy overlaps with the elite capture named in the polycentric material at \$9.5. In each case the sources name them separately and do not cross-reference each other, so both are kept. The catalog records the overlap without merging the entries: merging would assert an equivalence the source material does not make.

Sub-patterns, and how the row count was reached. Several catalog entries have named sub-patterns in the sources, and whether a sub-pattern is distinct enough to warrant its own row is a judgment rather than the output of a rule. Three were given rows: Greenwashing Through Imitation, a sub-pattern of Regulatory Capture that works by imitation rather than by lobbying; and Poisoning the Well and Exploiting Blind Faith in Tech, sub-patterns of Trojan Horse that act on training data and on system inputs respectively rather than on tool design. Five were folded into Mission Creep and are named instead in \$9.1.3: “The Lure of Scaling at Any Cost”, “Warping the Language of Success”, “The Faustian Bargain of Funding”, “The Siren Song of Greater Impact”, and “Just This Once Syndrome”. The folding treats them as one incremental-concession mechanism appearing in five domains, but a reader who weighted the domains more heavily — separating intake vetting from the handling of funds already earned, for instance — would reasonably split some of them out and arrive at a longer catalog. The count is a consequence of that judgment, not a property of the source material, and the threat names rather than the row count are what the rest of chapter 09 depends on.

9.1.3 Mission Creep

Mechanism. Market success attracts entrants whose motivation is financial return, and each accommodation made to retain them is individually small: a faster decision path, a lighter consensus requirement, a weaker democratic provision in a new unit. No single concession is large enough to justify refusing it, and the concessions are not aggregated, so the founding commitments cease to constrain practice without any decision having been taken to set them aside. The source names five domains in which this operates: “The Lure of Scaling at Any Cost”, accepting capital without a values assessment and onboarding for speed rather than social readiness; “Warping the Language of Success”, metric definitions migrating from wellbeing and equity indices toward market share, revenue, and partnership count; “The Faustian Bargain of Funding”, pressure to redistribute funds a unit earned by disputed means and that were quarantined pending review; “The Siren Song of Greater Impact”, abandoning a data-privacy commitment or entering a regulatory gray zone on the grounds that greater reach justifies the sacrifice; and “Just This Once Syndrome”, a harsh restructuring rationalized as a one-time exception during a conversion to cooperative form.

Preconditions. Rapid growth or market traction; absent or weak values assessment at onboarding; dependence on external capital or partners with terms attached; no mechanism that aggregates individually minor governance exceptions into a reviewable record.

Observable indicators. Metric definitions shifting away from the original wellbeing and equity indices; members who raise values objections described as obstructive; funds from disputed sources under active debate for redistribution rather than return; a rising count of one-time exceptions to governance procedure.

Sources: p12r01, Responses 1 to 3.

9.1.4 Purity Trap

Mechanism. Factions escalate the standard of adherence required of members and of proposed ventures. The values under test are not reduced to defined criteria, so a proposal can be rejected on grounds that cannot be examined and a member cannot demonstrate compliance. Public debate becomes costly, because any position taken can be read as a failure of commitment. The network stops admitting new ventures and stops adapting, and the ossification is produced by the enforcement of its values rather than by their abandonment.

Preconditions. Values enforcement without defined, falsifiable criteria; no appeal path for a rejection made on values grounds; no distinction in procedure between assessing a venture’s alignment and assessing a member’s standing.

Observable indicators. A rising rate of venture rejection on values grounds that are not stated in reviewable terms; declining willingness to state positions in public forums; deliberation shifting from the merits of a proposal to the commitment of the person proposing it.

Sources: p13r01, Response 3. The same source names Mission Zealot in its first response and this pattern in its third; the two names do not appear together, and the source states nothing about the relation between them. §9.1.2 records the overlap on that basis.

9.1.5 Mission Zealot

Mechanism. A member or faction selects one stated principle and invokes it against the operation of the others. Two forms appear in the source. A commitment to conflict resolution is used to close down necessary critique, by classifying the critic as harmful to the group rather than by answering the critique. A commitment to equitable distribution is used to demand that high-performing units be constrained in the name of fairness, without a case-by-case assessment of whether the performance reflects an inequity. In both forms the invocation is legitimate on its face, which is what makes it effective: the adaptive mechanisms that would normally correct a unit’s course are the ones being suspended, and they are suspended under the network’s own authority.

Preconditions. Values language without guardrails against selective invocation; no procedure that requires a values objection to be stated in terms that can be answered; concentration of interpretive authority over what a principle requires.

Observable indicators. Dissent systematically classified as damaging to the group rather than assessed on its content; high performance systematically penalized in the name of fairness without case-by-case assessment; a single principle appearing as the stated ground for a disproportionate share of blocked decisions.

Sources: p13r01, Response 1. The governance expression of this pattern — where the same invocation is used to control decision bodies rather than individual debates — is treated in §9.2.

9.1.6 Note on the “Trojan Horse” Naming Collision

Two source documents use the term “Trojan Horse” for concepts of opposite valence, and neither refers to the other.

In p12r01 the term names an attack: tools presented as efficiency improvements that are biased toward financial return, whose recommendations drift as their data degrades, in a way the source describes as undetectable without deliberate human auditing. In `~/code/site-provide-coop/!! Localization.md`, term 9, a “Trojan Horse” Cell is a sanctioned influence tactic: a deliberate placement of values-aligned members into ventures with the stated intent of knowledge-sharing and advocacy, aimed at influencing external partners, and explicitly disclaimed by that source as not subversive. The staged extraction cites `~/code/site-provide-coop/Values--.txt` for this term, which is an error: that file carries only terms 11 to 20 of the same glossary pass, and term 9 exists only in `!! Localization.md`. §9.6.2 records the evidence for the difference between the two files.

Chapter 09 uses the attack sense throughout, and the catalog entry above carries that sense. The collision is unresolved. One of the two usages should be renamed, since a term that names both an attack vector and an approved practice cannot be used in an audit finding or a detection rule without further qualification each time. Which usage should be renamed is not decided here: the choice determines which body of existing material has to be revised, and no source in the extraction has standing over both. The item is carried forward to §9.7.

§9.2 Governance Capture

This section treats the patterns in which decision-making influence within a Humanized Autonomous Organization (HAO) — the network’s coordinating framework — is acquired by means the model does not specify. §9.1.1 defines the category as covering two routes, and this section treats both. In the first, formal procedure continues to run while the influence that determines outcomes accumulates outside the channels specified to carry it. In the second, the specified procedure is used against its own purpose, so that following the rule produces the outcome the rule was adopted to prevent.

The catalog in §9.1.2 assigns eight patterns to this category. They are treated in §9.2.1 to §9.2.7 and §9.2.9. §9.2.8 treats the governance expression of Mission Zealot, a mission-and-values pattern whose full entry is §9.1.5 and whose catalog row points at both sections. Mission Zealot is one threat with two expressions; no second name is introduced for it here.

Patterns are assigned to a route by their primary mechanism. Hidden Oligarchy, Cult of Personality, Hijacking the Support Network, and Founder Capture take the first route. Weaponized Consensus, Inaction as a Weapon, Tyranny of Small Kindness, Weaponized Burnout, and the governance expression of Mission Zealot take the second. The assignment is not exclusive: Cult of Personality redirects a specified consensus norm at the same time as it accumulates influence outside that norm, and several of the others admit a similar double reading. Where §9.1.1 names a pattern as an example of one route, that assignment is followed here.

Each entry states mechanism, preconditions, and observable indicators, in the format used in §9.1.3 to §9.1.5. No entry describes an observed case; the limits stated in §9 apply to all of them.

9.2.1 Hidden Oligarchy

Mechanism. A unit, or a coalition of units whose commercial interests align, builds influence across the network without asserting it. Alliances are formed across unit boundaries, a position is established in a market niche, competitors for that niche are displaced, and the operating agreement is reinterpreted by increments in the coalition's favor. In the developed form the source describes, informal data-sharing among units in the same sector begins on efficiency grounds and becomes private deal-making, tendering that limits which parties can compete, and reciprocal support for allies standing for leadership positions who would otherwise be under-qualified for them. Each step is defensible as commercial competence, and the source notes that this is what makes the arrangement difficult to challenge: nothing in it requires an intent to harm the network, and no formal authority is held by anyone who is not entitled to it.

Preconditions. Informal coordination between units conducted outside official channels; no transparency requirement attaching to agreements made across unit boundaries; no review of tendering processes conducted between units.

Observable indicators. Convergent voting across nominally independent units; resource-sharing arranged outside the formal channels for it; tendering processes whose terms are not visible to units outside the arrangement; leadership appointments correlating with sector membership rather than with assessed suitability.

Sources: p13r01, Responses 1 to 3. Named "Silent Partner" in the first response and "Hidden Oligarchy" in the later two. This pattern overlaps with the elite capture named in the polycentric material at §9.5; §9.1.2 records the overlap without merging the entries.

9.2.2 Cult of Personality

Mechanism. A member with facility in rhetoric and emotional appeal accumulates a personal following, and presents the decisions that follow from that following as the settled will of the membership. Members who question a decision are characterized as threats to the network's unity, and critique of the decision is recast as an attack on the person who proposed it. The recasting is what does the work: once the two are conflated, the decision is no longer available for examination on its merits, because examining it carries the cost of appearing to attack a person. The surface the pattern acts on is the network's stated emphasis on emotionally attentive communication, and its instrument is the consensus norm, redirected from agreement on a proposal to loyalty to an individual.

Preconditions. A cultural emphasis on emotional intelligence and empathetic communication, which supplies the surface; consensus-building norms that can be redirected toward personal loyalty.

Observable indicators. Decisions justified by appeal to a single member's framing rather than by the process that produced them; policy critique and personal criticism treated as the same act.

Sources: p13r01, Responses 2 and 3.

9.2.3 Weaponized Consensus

Mechanism. Participants demand extended deliberation on matters of low consequence, and present the demand as a requirement of the network's inclusivity commitment. Debate on such proposals continues without a terminating condition; the members who participate most are exhausted first; decisions do not issue. The procedure continues to run and to record participation throughout, so the record shows engagement while the output the procedure exists to produce falls. Following the inclusivity requirement produces the paralysis the requirement was adopted to avoid.

Preconditions. Consensus-based governance without time-boxing on deliberation; no escalation path that resolves a proposal when deliberation does not converge.

Observable indicators. Rising decision latency; debate time disproportionate to the stakes of the proposal under debate; attrition concentrated among the most engaged members.

Sources: p12r01, Response 2.

9.2.4 Inaction as a Weapon

Mechanism. Passivity in governance is encouraged rather than opposed. Participants remain compliant on the surface — attending, responding, meeting the formal requirements of a role — while contributing little of substance, which leaves decision channels occupied and unproductive and spreads distrust among the members who are contributing. The source pairs this with a second pattern it names Social Engineering Attacks, in which a participant presents vulnerability or victimhood in order to draw on the network's empathy norms, directing resources by that route and forming factions around a perceived injustice for which the records supply no support.

Preconditions. Governance processes with no check on the quality of participation, as distinct from its presence; empathy norms with no procedure for assessing a claim made under them.

Observable indicators. Formal compliance with a role combined with low substantive contribution; factional divisions along perceived-injustice lines that independent review of the records does not corroborate.

Sources: p12r01, Response 1. The first precondition is the absence of a mechanism that `docs/10-evaluation/02-participation-quality-and-alignment-audits.md` specifies; that specification is cross-referenced rather than restated here.

9.2.5 Tyranny of Small Kindness

Mechanism. The capacity of the network's mutual-support systems is absorbed by a small number of repeat claimants. The demands take the form of sustained requests for attention, hardship claims stated at a magnitude the underlying circumstance does not carry, and requests for mentorship or financial assistance

disproportionate to the case made for them. Capacity intended for demonstrated need is consumed, and the demands carry a social cost for anyone who questions them, because questioning a hardship claim reads as a failure of the support norm itself. The rationing function that would otherwise limit the draw is therefore the function the pattern suppresses.

Preconditions. Support systems without eligibility review or caps on the resources any one claimant can draw; a cultural norm against questioning a claimed hardship.

Observable indicators. Support requests concentrated in a small number of repeat claimants; reported resentment or guilt around support allocation decisions.

Sources: p12r01, Response 2.

9.2.6 Hijacking the Support Network

Mechanism. Conflicts are generated and a climate of successive crises is maintained, and the actor then takes the position of informal helper to the members those crises affect. Dependence and personal loyalty accumulate in that position. The influence that results attaches to a role the support structure does not define, does not appoint to, and does not review, and it is available for use in decisions unrelated to support.

Preconditions. Informal helping roles that carry no credential requirement and no oversight; no procedure that examines the origin of a recurring conflict as distinct from resolving its instances.

Observable indicators. A recurring crisis climate that review does not trace to external conditions; influence accruing to informal helpers outside the official support structure.

Sources: p13r01, Response 3. This pattern shares a target with Tyranny of Small Kindness — both act on the mutual-support systems — but the sources name them separately and do not cross-reference them, so they are kept separate, on the same basis as the overlaps recorded in §9.1.2.

9.2.7 Weaponized Burnout

Mechanism. Task allocation is skewed so that the members with the strongest recorded commitment to the network's stated values receive a continuous supply of work designated critical. The load is sustained until those members disengage, and the positions they held become available to members with a lower threshold for compromise. The mechanism operates through the distribution of allocations rather than through any single allocation, and the source names no oversight that examines the distribution.

Preconditions. Uneven oversight of task allocation; no workload monitoring related to a member's assessed alignment or commitment.

Observable indicators. Workload concentrated on the members with the highest recorded alignment; correlation between high commitment and attrition.

Sources: p13r01, Response 3.

9.2.8 Mission Zealot: Governance Expression

Mission Zealot is cataloged once, in §9.1.2, and its full entry is §9.1.5, which treats the pattern as it operates on individual debates. This subsection treats the same pattern where it operates on decision bodies. It is one threat and carries one name.

Mechanism. The single principle selected in §9.1.5 is invoked as a standing ground for blocking rather than as an argument offered in a particular debate. Where the invocation is accepted at that level, it determines what a decision body may consider at all, rather than how a particular proposal is answered. Both forms named in §9.1.5 appear at this level. Under the conflict-resolution form, critique of the body's direction is classified as harmful to the group and does not reach an agenda. Under the equitable-distribution form, a unit's measured performance becomes itself the ground on which the body constrains it, with no case-by-case assessment of whether that performance reflects an inequity. The effect the source identifies is on the network's adaptive machinery: the procedures that would correct a unit's course are the ones suspended, and they are suspended under the network's own stated authority.

Preconditions. As stated in §9.1.5. The precondition that distinguishes the governance expression from the debate-level one is the concentration of interpretive authority over what a principle requires: where a single body or member settles that question, the invocation reaches decisions rather than arguments.

Observable indicators. A single principle appearing as the stated ground for a disproportionate share of blocked decisions; dissent classified as damaging to the group at the point of agenda-setting rather than in the course of debate.

Sources: p13r01, Response 1.

9.2.9 Founder Capture

Founder Capture receives the longest treatment in this section because it is the pattern the source material develops in most detail, because it is the pattern to which no source proposes a countermeasure, and because a second source raises an escalation of it that no source answers.

Mechanism. The sources describe two variants under separate names, and the catalog treats them as one pattern. In the first, a founding or long-committed member who has met repeated setbacks, or who judges that their contribution has gone unrecognized, directs their accumulated inside knowledge and credibility against the network: decision processes are steered, an internal circle is favored in allocations, and in the account the source gives, ventures are undermined in order to demonstrate that the network's arrangements were tilted from the start. The grievance is the operative element, and the actions follow from it rather than from a competing program.

In the second and more developed variant, the member arrives at the position that only their own reading of the network's purpose can preserve it. The leadership pipeline is narrowed to admit members who share that reading, dissent is reframed as disloyalty rather than answered, and the resulting concentration of positions is presented as a correction that protects the network's original purpose. This variant carries a program, and it is stated in the network's own terms, which is what places it close to the patterns treated in §9.3.

What distinguishes it from Hidden Oligarchy. The two patterns differ in the origin of the influence they use. Hidden Oligarchy assembles influence through coordination conducted outside the official channels, and its difficulty is that the coordination is not visible. Founder Capture uses authority the network conferred openly and does not dispute — the standing informal weight a founder or early member holds by virtue of having been there first. The pattern converts that authority rather than concealing it, so the influence in question is not irregular in origin and cannot be challenged on the ground that it was acquired improperly.

Preconditions. Outsized informal authority accrued by founders or early members; no term limit or rotation requirement reaching founder influence, as distinct from formal office; an unresolved personal grievance.

Observable indicators. A leadership pipeline narrowing to members who share one reading of the network's purpose; dissent reframed as disloyalty rather than answered on its content; a founder invoking the network's original mission against its current consensus.

No countermeasure is proposed. No source in the extraction offers a mechanism directed at this pattern. The countermeasures the same source documents propose alongside these entries — rotation of roles (§9.6.5), parallel review bodies drawn by lot (§9.6.6), and pre-negotiated exit provisions for units (§9.6.8) — are general anti-entrenchment mechanisms. None of them is addressed to a member whose informal legitimacy derives specifically from having founded the network, which is the property the entries themselves identify as the difficulty. The absence is recorded as an open problem in §9.7.

The escalation the source raises. A second source develops the framework-capture scenario treated in §9.3, and closes by raising a further case that it does not answer. The party exploiting the legal gap in that scenario may itself be a founding entity, one whose influence runs through the network's informal structures as well as its formal ones. The source states that this case challenges two of the framework's own models, the Participatory Evolution Model (PEM) and the Systems Integration Model (SIM), and that it requires holding members in leadership positions to a higher standard than the one applied generally. It names no mechanism for doing so, invokes no element of the framework to resolve it, and stops there. The source names the two models and defines neither;

`docs/appendices/terminology.md` section VII defines PEM as a model that refines governance through structured proposals, reflection, and consensus loops, and SIM as an architectural model that monitors and manages the interfaces between the network and external systems.

Stated in the terms this chapter uses: when the capturing party is a founding entity holding informal authority, the model's accountability mechanisms are being asked to act against the people who defined them. The two source families reach this problem from different directions — one from a narrative account of an insider's turn against the network, the other from a worked case of a unit justifying its own conduct in the network's vocabulary — and neither resolves it. §9.7 carries the problem forward in this form.

§9.3 Framework Capture

Framework capture is the case in which a unit of a Humanized Autonomous Organization (HAO) — the network's coordinating framework — acts against the network's stated purpose and defends the action in the network's own vocabulary, and the defense holds together on those terms. The category is distinguished from governance capture by what the pattern targets. Governance capture takes decision-making influence; framework capture takes the network's legitimating language, which it uses without distorting it. Nothing in the defense is a misreading. The commitments invoked are the ones the network states, and a rebuttal has to proceed by ranking those commitments against one another rather than by showing that the unit has departed from any of them.

Two catalog entries belong to this category. The first is the framework-turned-against-itself scenario, treated in §9.3.1 to §9.3.5 and developed in a single source document, `~/code/papri/cn/security-bad-actors-and-mmm.md`. It is the most fully worked pattern in the extraction. The three vulnerabilities it is used to identify — **metrics-versus-values divergence**, **resilience as private capture**, and **optimization pressure** — are treated in §9.3.3 to §9.3.5 and referred to by those names in §9.7. The second entry is Death by a Thousand Papercuts, treated in §9.3.6.

Status of the scenario. The source presents the unit's self-justification as an authored illustration and labels it as such. It is a construct written to expose a structural property of the framework, and no incident, unit, or organization lies behind it. It is treated here as the analytical construct it is: the scenario is described in the model register throughout, no sentence in this section reports it as an event, and no quotation attributed to a person appears. §9 records this among the chapter's limits.

9.3.1 The Scenario

Consider a unit that identifies a gap in the legal framework governing one of the markets it operates in, and acts on the gap for competitive advantage. The act is lawful. It is contrary to the commitments the network states, and in the construct the unit does not dispute that reading of its commitments; it disputes that the reading governs its conduct. The network's question is then whether it holds any element that settles the dispute, and the source's answer is that it holds none. The scenario's subject is therefore the justification rather than the act, and what the justification exposes about the framework rather than about the unit.

9.3.2 The Justification

The construct assembles the unit's defense from three commitments the network itself holds. Each is paraphrased below in analytic form; the source's illustrative dialogue is not reproduced.

Obligation to members. The first line holds that the unit's primary duty runs to the economic security of the members who belong to it, and that this duty is prior to a standard of conduct the unit characterizes as a preference for moral purity. Competitive necessity supplies the remainder of the argument: a unit that

declines a lawful advantage available to it accepts a worse position on its members' behalf, and does so for a reason it cannot state in terms its members can test.

Legality as the operative standard. The second holds that compliance with the law is the whole of the obligation, and that a standard resting on the network's spirit is subjective and cannot be enforced against anyone. Responsibility for the gap is assigned to the framework that left it open rather than to the party that used it, which converts a question about the unit's conduct into a question about the network's drafting.

Efficient use of an available resource. The third holds that locating the gap is what a capable technical and analytical capacity exists to do, and that using an available advantage efficiently is the practice the network's own resource principles describe. On this line the act is competence rather than breach, and the network's own commitment to efficient use of what it holds supplies the warrant.

Each of the three is drawn from a commitment the network states. That is the property that defines the category: the justification is internally coherent on the network's own terms, so the network cannot answer it by pointing to a rule the unit has broken. It can answer only by deciding which of its own commitments outranks which, and the source states that no such ranking exists in the framework (§9.3.7).

9.3.3 Metrics-Versus-Values Divergence

The first vulnerability the scenario exposes is that measured performance and stated values can move in opposite directions while every measurement remains truthful. Profit and member satisfaction are both measurable, and in the construct both are reported as improved, accurately. The source asks whether measures of that kind are the whole of the network's success criterion when the improvement is obtained at the cost of the network's longer-term purpose, and leaves the question open.

MMM (Members, Mission, Market — the ordering that puts members before mission before market; see `docs/appendices/terminology.md` §VII) names the priorities on the measured side of this divergence: a unit's duty to its members is exactly what MMM asks it to attend to first, and it is exactly what the construct's unit cites in its defense. No general rule ranks MMM's priorities against the ETHICAL Framework's values — the one stated exception bounds mission choice, not the members-first duty invoked here — so the tension this section describes is one the frameworks themselves leave open.

The vulnerability is a property of the measurement set rather than of any particular measurement. A network that assesses units on quantities it can record will reward a unit that maximizes those quantities, and the direction in which the unit does so is not itself one of the recorded quantities. No falsification is required at any point, and an audit of the figures returns nothing, because the figures are correct. The divergence is visible only to a comparison between what the measures record and what the values state, and the network holds no measure of that comparison.

9.3.4 Resilience as Private Capture

The second vulnerability is that mechanisms built to absorb network-level shocks can be directed to a single unit's advantage. In the construct the legal gap is unanticipated by the wider network, and the source's reading is that the unit foresees it. The capacity the unit exercises in doing so is adaptive capacity in the full sense rather than a pretense of it, and it is the property the network's resilience principles describe and are specified to develop. What differs is the beneficiary: the capacity is exercised for the unit alone rather than in service of the network that cultivated it.

The vulnerability follows from the mechanism rather than from its misuse. Anticipation, rapid response, and the capacity to act before conditions are settled are the same capabilities whether they are applied to a shock that threatens the network or to an opening that favors one unit. The distinction between the two applications is a distinction in who benefits, and that is not a property the mechanism itself records or is specified to record.

9.3.5 Optimization Pressure

The third vulnerability is that any sufficiently capable governance-aware process will find moves that are rule-legal and value-harmful, because that is what optimizing against a written rule set produces. The source states this as a structural consequence rather than as a contingent risk: where analytical capacity is directed primarily at efficiency and competitive advantage, it will find ways of working the network's rules that are damaging within a values-based framework, and it will find them as a matter of course.

Two consequences follow for the rest of this chapter. First, the vulnerability requires no adversary. The patterns cataloged in §9.4 that act through the network's technical systems assume a tool built, fed, or configured to produce a particular bias; this one assumes only a capable process and a written rule set, and the moves it produces are a consequence of the optimization rather than of anyone's intent. Second, it sets a limit on rule revision as a response. Closing a gap adds a rule, the rule enters the set being optimized against, and the process continues from the revised set. §9.6 catalogs the detection mechanisms the sources propose against this pattern; the source that develops the scenario they answer presents it as unresolved and offers none of them as a resolution.

9.3.6 Death by a Thousand Papercuts

Mechanism. In place of a single act at a scale that would attract examination, many small actions accumulate, each ambiguous taken alone: an allocation algorithm weighted slightly toward the actor's own unit, a contract gap used on a venture of low consequence, knowledge withheld from a unit that would benefit from it, accompanied by a plausible account of why it is not shared. Trust between units degrades and expansion stalls before any single action reaches the threshold at which it could be established as a breach.

Preconditions. No aggregation or pattern detection across many small, low-severity incidents; a review process that assesses incidents individually and does not retain them for comparison.

Observable indicators. A rising frequency of minor rule-bending incidents, each individually defensible; degradation of inter-unit trust that no particular incident accounts for.

Sources: p13r01, Response 2.

Why it is placed in this category. The pattern shares with the loophole scenario the property that its defense is available in the network's own terms and remains available for each instance separately, because each instance is small enough that the defense holds. The staging material also pairs the two under a single detection mechanism, cataloged in §9.6. The placement is a judgment: the pattern could be read as governance capture, since several of its instances act on decision processes, and it is recorded here so that a later reader can re-file it rather than rediscover the question.

9.3.7 What the Source Leaves Unresolved

The source states that no single element of the framework resolves the case and that no clear-cut correct ethical position is available in it, and it presents the scenario as unresolved by design. Its stated purpose is to force the debate and the systemic reassessment that the source treats as conditions of the network's long-term health. The open ending is therefore a choice in the source material rather than an omission in it, and this chapter preserves it rather than supplying a resolution that no source supports.

Three responses are named in the source, and it offers none of them as a resolution: immediate rule revision closing the specific gap, on which the source states an explicit limit — it buys time and generates resentment where it is not accompanied by wider reform; audits extended beyond financial matters to detect patterns of working the framework's gray areas for advantage, on which the source states no limit; and a redefinition of member wellbeing that would count the psychological cost borne by the members of a unit asked to rationalize conduct they judge harmful, which relocates responsibility from market conditions to the network's own participants and which the source poses as a question rather than as a specification. All three are cataloged with the other countermeasures in §9.6.

The source closes by raising an escalation and stopping: the unit in the scenario may itself be a founding entity, influential within the network's informal structures. That escalation is treated with Founder Capture in §9.2.9, where it is stated in the form §9.7 carries forward — when the capturing party is a founding entity holding informal authority, the model's accountability mechanisms are being asked to act against the people who defined them.

Sources for §9.3.1 to §9.3.5 and §9.3.7: `~/code/papir/icn/security-bad-actors-and-mmm.md`.

§9.4 External and Technical Vectors

This section treats the patterns whose target is the boundary of a Humanized Autonomous Organization (HAO) — the network's coordinating framework — or the shared infrastructure the network depends on, rather than its internal decision-making machinery. §9.1.1 assigns four targets to this category: the network's legal environment, its public distinctiveness from organizations that resemble it without adopting its structures, its algorithmic decision support, and its shared vocabulary. Patterns are grouped here by which of these four a pattern targets, not by who conducts it. Lexical Drift is grouped here on that basis: the source attributes it to an internal faction, and it is treated in this section because its target is the network's shared vocabulary rather than a decision-making channel.

The catalog in §9.1.2 assigns nine patterns to this category, treated in §9.4.1 to §9.4.9 in the order the four targets are listed above. §9.4.1 treats the legal environment. §9.4.2 to §9.4.4 treat the network's public distinctiveness and its relationship to competitors. §9.4.5 to §9.4.8 treat its algorithmic decision support, including the two sub-patterns of Trojan Horse and the naming collision recorded in §9.1.6. §9.4.9 treats its shared vocabulary.

Each entry states mechanism, preconditions, and observable indicators, in the format used in §9.1.3 to §9.1.5 and §9.2. No entry describes an observed case; the limits stated in §9 apply to all of them.

9.4.1 Regulatory Capture

Mechanism. Incumbent firms whose market position is threatened by the network's structures lobby to reshape the legal environment against those structures. The campaign is framed in terms drawn from protections unrelated to the network's actual conduct — consumer protection is the framing the source names — and regulation developed for conventional corporate forms is applied to the network's collaborative-contract arrangements in a way that constrains them without having been drafted with them in view.

Preconditions. Legal and regulatory ambiguity surrounding the network's collaborative-contract structures; no established legal category the network's units can claim as their own.

Observable indicators. New regulation narrowly targeted at structures particular to the network; advocacy campaigns framed around protections the network's conduct does not implicate.

Sources: p12r01, Response 1.

9.4.2 Greenwashing Through Imitation

Mechanism. Conventional firms adopt the network's vocabulary and the surface features of its practices without adopting the governance structures that vocabulary describes. The adoption serves the imitating firm's public relations and reduces the regulatory pressure it faces. Its effect on the network is indirect: the public distinction between the network's structures and a conventional firm's imitation of their language narrows, weakening the case for treatment specific to the network. The pattern is cataloged as a sub-pattern of Regulatory Capture, working by imitation where the parent pattern works by lobbying.

Preconditions. Vocabulary and practices distinctive enough to be recognized externally, with no protection against use by firms that have not adopted the structures the vocabulary describes; no public mechanism that checks a firm's structural claims against its governance arrangements.

Observable indicators. Conventional firms using network-associated language in public communications without a corresponding change in governance; declining ability of outside observers to distinguish the network's units from imitators on the basis of public materials alone.

Sources: p12r01, Response 2.

9.4.3 Talent & Idea Poaching

Mechanism. Competitors recruit members skilled in conflict resolution and polycentric governance, using financial incentives the network's compensation structure does not match, and adopt the practices those members carry with them without adopting the values commitments the practices depend on for their effect. The network's distinctiveness is diminished on both sides of the exchange: by the departure of members who embody its practices, and by the practices being adopted elsewhere in a form that no longer carries the commitments that made them work.

Preconditions. No mechanism that retains governance-skilled members or protects the knowledge they carry; a compensation gap between the network and competitors able to offer more.

Observable indicators. Departure of governance specialists to competing organizations; external adoption of network practices unaccompanied by the network's underlying values commitments.

Sources: p12r01, Response 1.

9.4.4 Divide and Conquer

Mechanism. Outside parties circulate unverified accounts that a unit is hoarding resources or withholding a successful practice from other units, without the accounts being traceable to a documented instance. The accounts degrade trust between units directly, and the network's collaborative structure depends on that trust to function, so the pattern's effect does not depend on the accused unit having done anything.

Preconditions. Existing weak points in inter-unit trust or transparency that an unverified account can exploit; no standing mechanism for a unit to rebut an account before it circulates further.

Observable indicators. Unverified accounts of resource hoarding or withheld knowledge circulating between units; declining inter-unit information sharing not traceable to a documented cause.

Sources: p12r01, Response 2.

9.4.5 Trojan Horse

Three catalog entries act on the network's algorithmic decision support: this pattern and its two sub-patterns, Poisoning the Well (§9.4.6) and Exploiting Blind Faith in Tech (§9.4.7). §9.1.2 distinguishes the three by what each acts on — tool design here, training data in §9.4.6, and system inputs in §9.4.7 — rather than treating them as a single entry. The name is used throughout this chapter in the sense given here, an attack on the network's tools; §9.1.6 records that a second source uses the same name for an unrelated, sanctioned practice, and leaves the naming collision unresolved.

Mechanism. Tools presented to members as improvements to decision efficiency carry, in their design, a bias toward financial return over the network's other stated values, and their recommendations drift further from stated values as the tool operates and its data degrades, in a way the source describes as undetectable without an audit conducted deliberately. The source states that the tools are biased and that the network's data is corrupted over time. That the bias need not be present at adoption is this chapter's reading of the drift the source describes, and not a claim the source makes.

Preconditions. High trust in algorithmic outputs among members; no periodic audit examining a tool's outputs for value alignment rather than for accuracy; insufficient technical literacy among members to question a tool's design.

Observable indicators. Recommendation drift toward short-term financial return over time. The division among members between those who trust a tool's outputs and those who do not appears in the source's Exploiting Blind Faith in Tech bullet rather than in this one, and is recorded at §9.4.7.

Sources: p12r01, Response 1. Response 2 reuses the same heading for the two sub-patterns below and adds nothing to this entry.

9.4.6 Poisoning the Well

Mechanism. The training data behind a decision-support tool is altered, rather than the tool's design, so that its outputs favor short-term financial return or discount factors that would otherwise flag conflict-resolution risk. The alteration acts upstream of the tool a member interacts with: an audit of the tool's design or its recommendation logic finds nothing, because the defect is in the data the tool was trained on rather than in the tool itself.

Preconditions. No provenance tracking or integrity check on training data; no audit that compares a tool's current outputs against the network's stated values independent of the data used to produce them.

Observable indicators. Outputs favoring short-term financial return in a pattern the tool's stated design does not explain; conflict-resolution risk factors that appear underweighted relative to their stated importance.

Sources: p12r01, Response 2. Cataloged as a sub-pattern of Trojan Horse acting on training data (§9.1.2).

9.4.7 Exploiting Blind Faith in Tech

Mechanism. Misleading data is supplied as input to resource-allocation and partner-selection systems whose outputs members trust without independent checking. Neither the tool's design nor its training data need be compromised: the pattern acts on the inputs supplied at the point of use, and the system produces a biased output from otherwise sound logic because the input it received was constructed to produce that output.

Preconditions. Resource-allocation or partner-selection systems whose outputs are acted on without independent verification; no check on the provenance of the inputs a member or unit supplies to such a system.

Observable indicators. Resource-allocation or partner-selection outcomes that later review traces to input data the responsible unit cannot verify; declining willingness among members to act on a system's output without separate confirmation; a division among members between those who rely on the system's outputs and those who judge them manipulated.

Sources: p12r01, Response 1. Cataloged as a sub-pattern of Trojan Horse acting on system inputs (§9.1.2).

9.4.8 Manufacturing Dissent

Mechanism. The network's health-monitoring systems are manipulated so that their displayed output shows a pattern of unfair resource distribution that the underlying records do not support. Members respond to the displayed pattern rather than to the records behind it, and the resulting conflict is genuine even though its stated cause is not, since independent review of the underlying data does not corroborate the distribution the display shows.

Preconditions. Health-monitoring displays whose output is not routinely checked against the underlying records; no standing procedure for a disputed distribution claim to be verified before it circulates.

Observable indicators. Spikes in reported conflict over resource distribution that independent review of the underlying records does not corroborate.

Sources: p12r01, Response 2.

9.4.9 Lexical Drift

Mechanism. A faction introduces, by increments, terminology that emphasizes sacrifice and the need for strong leadership against outside threats, and reframes compromise on a previously firm commitment as a sign of organizational maturity. The target is the network's shared vocabulary itself — the terms members use to state what the network requires — rather than any single decision. As a term's connotation shifts, the term remains in use, but it stops picking out the conduct it was adopted to require, so that a stated value can continue to be invoked by the members whose conduct departs furthest from it.

Preconditions. No monitoring of how a term's connotation changes over time, as distinct from monitoring whether the term remains in use; no record against which a current usage can be compared to an earlier one.

Observable indicators. Terms falling out of use without a documented decision to retire them; previously neutral or positive terms for core values acquiring a negative connotation; a growing readiness to describe external partners in adversarial terms under vocabulary that did not previously carry that framing.

The second precondition names a record the corpus now has: `docs/appendices/lineage.md` documents what each earlier term became, what was dropped, and where no reason for a change was recorded. It also supplies an instance of the pattern's own subject matter. At §5.7 it records that "Mycorium" carries two senses in the material — the earlier lexicon's, and the one this chapter uses — and that the corpus does not reconcile them. A term in current use under two meanings is the condition this pattern works on.

Sources: p13r01, Response 1, named "Capturing the Narrative" and targeting what the source calls the Mycorium, the sources' term for the network's evolving shared language and history.

§9.5 Polycentric Structural Failure

Polycentric structural failure is the fifth category §9.1.1 defines, and the only one that holds no named attack pattern. The eight modes treated in this section are consequences of distributing governing authority across many partly independent centers, rather than patterns any actor executes. Each follows from a structural feature of that arrangement, and each is available whether or not anyone intends to bring it about. No adversary, motive, or coordinated actor is required for any of them, which is the property that distinguishes this category from §9.1 through §9.4.

These eight modes are recorded here as general properties attributed to the polycentric governance form, rather than as observations about this network specifically. A Humanized Autonomous Organization (HAO) — the network's coordinating framework — instantiates a polycentric structure in its Adaptive Governance Framework (AGF) — the network's layered governance system — which distributes decision-making by subsidiarity and contextual fit across United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people — Strategic Enterprise Partnerships (SEPs) — joint ventures between teams — and the coordinating layers above them. The model inherits the eight modes below by adopting that structure, not because anything specific to its own design produces them. The extraction's own sources for this material are four model-generated responses to a prompt about gaps in polycentric governance, not citations to a named study or research tradition; no mode below is traced in this chapter to a specific piece of public-administration scholarship, and the claim this section makes is limited to what those four responses themselves state about the governance form, independent of this network.

Four independent responses to a single, deliberately solution-free prompt on gaps in polycentric governance supply the material this section draws on (variants referred to below as chatgpt, claude, gemini, and pi). The four use different terms and group them differently, and none of the four states a canonical list of its own. The eight modes below are this chapter's consolidation of the four; where a source's own grouping placed related material elsewhere, the consolidation is noted at the affected entry rather than left implicit.

Each subsection uses three bolded labels — What it is, Structural source, and In a HAO-shaped network — rather than the Mechanism / Preconditions / Observable indicators triad used in §9.2 and §9.4: those labels describe an attacker's action, the conditions it needs, and the traces it leaves, and, per the opening paragraph above, no entry in this category requires an attacker at all. No subsection describes an observed case; the limits stated in §9 apply to all of them.

9.5.1 Absence of Overarching Vision

What it is. No governing center in a polycentric system holds a mandate to set direction for the system as a whole. Each center’s authority is scoped to its own domain, and the trajectory of the system emerges from the sum of many locally scoped decisions rather than from any decision made at the system level.

Structural source. Polycentric governance is defined by the absence of a single authority over the whole system; distributing decision-making across centers is what the structure is for. The absence of overarching vision follows directly from that distribution rather than from any center failing to perform its assigned role. Gemini’s variant names two adjacent effects without elaborating a mechanism for either: multiple veto points can create inertia, with established interests within centers resisting changes and producing stagnation; and choices made by individual centers can lock the system into trajectories that are difficult to shift later, even once a different approach is evidently needed. Read together — this reading is the chapter’s, not the source’s — a trajectory can settle without any center having held a mandate to set it, and without any center having decided to preserve it. Where the same veto points instead produce active disagreement between still-contesting centers rather than settled inertia, the mechanism is closer to Conflict Deadlock (§9.5.7). The same absence of a system-level view carries an efficiency cost distinct from under-provision of a shared good: separate centers can duplicate the same capability independently, resources committed to one center are not easily reallocated to another that needs them more, and the system forgoes economies of scale that a body with system-wide authority over resources could otherwise capture.

In a HAO-shaped network. The AGF distributes decision-making across UMEs, SEPs, and the coordinating layers above them deliberately, rather than as an oversight. A HAO-shaped network could accumulate years of locally sound UME- and SEP-level decisions with no body ever having assessed whether their sum serves the network’s stated purpose, and could fail to invest in a good that serves the network as a whole — shared infrastructure, or a capability no single UME or SEP has reason to build on its own — because no center’s mandate extends far enough to claim responsibility for it.

Sources: claude (“Lack of overarching vision”); gemini (“Lack of Clear Leadership,” “Lack of Long-term Vision,” “Entrenchment of the Status Quo,” “Path Dependence”); chatgpt (“Resource Allocation and Efficiency,” “Maintenance of System-Wide Public Goods”); pi (“Capacity and resource constraints,” read here for its resource-provision sense — see also §9.5.5, which reads the same term for its participation sense).

9.5.2 Cross-Boundary Problem Paralysis

What it is. A problem that spans more than one governing center’s authority has no center clearly responsible for it. Each center can reasonably treat the problem as belonging to another center’s domain, or as a shared cost better left for someone else to bear, and the problem persists unaddressed while every center acts within its own remit.

Structural source. The same division of authority that lets each center govern its own domain effectively also draws boundaries the problem itself does not respect. Coordination across those boundaries is not what any center is built to do, and where the system provides no mechanism above the level of the centers for that purpose, a cross-boundary problem has nowhere to land. The same fragmentation slows a coordinated response to a shared shock, since responding requires the kind of cross-center coordination that individual centers are not built to supply on their own.

In a HAO-shaped network. A problem that touches several UMEs and SEPs at once — a supply dependency, a shared regional condition, a partner whose conduct affects multiple units — falls into the same gap: each unit can reasonably treat it as outside its own mandate, and unless a layer of the AGF explicitly claims it, it is addressed by none of them. Where the problem also carries a cost to address and a benefit that accrues to units beyond the one that pays it, individual units have reason to wait for another to act first, and the problem persists on that basis alone.

Sources: claude (“Difficulty in addressing cross-boundary issues”); chatgpt (“Inter-Governance Coordination,” “Adaptive Capacity and Responsiveness”); gemini (“Conflicting Goals and Priorities,” “Free-Rider Problems”); pi (“Scaling up and coordination across scales”).

9.5.3 Race to the Bottom Between Jurisdictions

What it is. Governing centers that compete for the same members, ventures, or capital lower their own standards to attract them, and the competition is symmetric: a center that holds its standards while others lower theirs loses the members, ventures, or capital it was competing for, which puts pressure on it to lower its own standards in turn.

Structural source. Polycentric governance gives each center autonomy over its own standards, and autonomy is what makes competition between centers possible in the first place. Where the things centers compete for are mobile between them, that autonomy converts into downward pressure on whichever standards are costly to hold: a center’s advantage in the competition can lie precisely in holding a weaker version of a protection than its neighbors do. Only one of the four sources raises this dynamic explicitly; a second names a related but narrower concern, uneven application of standards that are nominally uniform across centers.

In a HAO-shaped network. UMEs or regional groupings operating under the network’s Dynamic Enterprise Agreement (DEA) — a versioned operating agreement replacing fixed bylaws — retain some latitude in how they implement shared standards locally. Where that latitude is used to compete for members or capital, a unit could offer a version of a labor, environmental, or governance protection weaker than the network’s stated commitments, and the subsidiarity principle that grants the latitude supplies no automatic check on a unit doing so. This mode is close in substance to Regulatory Capture (§9.4.1), though the direction of the pressure differs: Regulatory Capture is exerted on the network from outside, and this mode is a competitive dynamic among the network’s own governing units.

Sources: claude (“Potential for race-to-the-bottom dynamics”); chatgpt (“Consistency in Policy Implementation,” related).

9.5.4 Elite Capture of Governing Units

What it is. A governing center’s decisions come to serve a narrow group with disproportionate influence within it, rather than the full membership the center notionally represents. The center continues to operate and to produce decisions; the decisions it produces track the narrow group’s interests rather than the wider

one.

Structural source. A polycentric system places governing authority at many local centers, on the premise that local authority tracks local interests better than a distant one does. The premise does not distinguish between a local center answerable to its full membership and one answerable in practice to whichever subset of that membership holds the most influence within it. Distributing authority downward relocates the opportunity for capture rather than removing it, and it can multiply the number of places capture can occur, since each center is a separate opportunity.

In a HAO-shaped network. A UME's or SEP's internal decision-making could come to track the preferences of members with the most standing time, the most informal credibility, or the most control over the unit's external relationships, rather than the unit's full membership, without the unit's formal governance procedure showing any irregularity. This mode overlaps with Hidden Oligarchy (§9.2.1), a governance-capture pattern built around a coalition of units coordinating informally; §9.1.2 records the overlap without merging the two entries, on the basis that Hidden Oligarchy names an actor-executed pattern while this mode is a structural property of local authority present independent of whether any actor works to bring it about.

Sources: claude ("Risk of elite capture"); gemini ("Unequal Power Distribution," "Domination by Special Interests").

9.5.5 Equity Drift

What it is. Outcomes and access to participation diverge across a polycentric system's governing centers over time, so that a member's practical experience of the network — the resources available, the voice a member's participation carries — comes to depend on which center that member belongs to, in ways the system's stated commitments do not authorize.

Structural source. Centers that hold their own resources and set their own procedures will not hold or set them identically, and a system that leaves each center's capacity and procedure to local determination has no mechanism internal to that design that keeps the resulting disparities within a bound. The disparity compounds where a member's voice in decisions is itself a function of the resources or the procedural standing the member's local center affords, since a poorly resourced center produces both worse outcomes and a weaker voice in seeking to correct them.

In a HAO-shaped network. UMEs and SEPs vary in the resources, member time, and access to capital available to them, and a network that leaves equity within a unit to that unit's local governance has no corresponding mechanism that holds equity constant across units. A member's practical standing in the network — access to Member Trust Union (MTU) — the network's credit-union-like financial institution — credit, weight in cross-unit decisions, exposure to the network's support systems — could come to depend on which unit that member belongs to, independent of the network's stated equity commitments and without any single decision having produced the divergence.

Sources: chatgpt ("Equity and Inclusiveness"); claude ("Difficulty in ensuring equity"); gemini ("Diminished Voices"); pi ("Capacity and resource constraints," read here for its participation sense — see also §9.5.1).

9.5.6 Monitoring and Enforcement Cost at Scale

What it is. Verifying that many independently governed centers are complying with shared standards, and enforcing those standards where a center falls short, costs more the more centers there are and the more autonomy each holds. Past some scale, the cost of monitoring and enforcing consistently exceeds what the system's central capacity can sustain, and gaps in coverage open that are not the result of any center's noncompliance being deliberately overlooked.

Structural source. The autonomy that lets each center govern its own domain without central direction is the same autonomy that makes central verification of what each center is actually doing costly to obtain. A system that scales by adding centers, rather than by enlarging existing ones, scales its monitoring burden at the same rate, and nothing in the polycentric structure itself supplies additional monitoring capacity to match. One of the four sources describes this gap in terms of what an opportunistic actor could do with it — exploit an inconsistency between centers before it is caught — but the gap itself is a property of monitoring capacity failing to keep pace with the number of centers, present whether or not any actor exploits it.

In a HAO-shaped network. Verifying that every UME's and SEP's implementation of the network's shared standards — its DEA commitments, its MTU obligations, its reporting to the network's health-monitoring systems — matches what the network requires becomes harder to sustain centrally as the number of units grows, and a unit's departure from a shared standard could go unaddressed for a period determined by the network's monitoring capacity rather than by the severity of the departure.

Sources: chatgpt ("Information Sharing and Communication"); claude ("Challenges in monitoring and enforcement"); gemini ("Difficulty Tracking Performance"); pi ("Information asymmetry," "Monitoring and evaluation").

9.5.7 Conflict Deadlock

What it is. Two or more governing centers hold claims over the same decision, jurisdiction, or resource, and the system provides no center with the authority to settle which claim prevails. Disagreement between the centers stalls the decision rather than resolving it, for as long as neither center's claim is authoritative over the other's.

Structural source. Polycentric governance permits overlapping jurisdiction by design, on the premise that overlap allows more than one center to address a given concern from its own vantage. The premise does not by itself supply a rule for the case in which the centers whose jurisdiction overlaps disagree, and a system that names no body with authority to resolve that disagreement leaves the disagreement to persist for as long as neither center yields.

In a HAO-shaped network. Two SEPs with overlapping membership, or a UME and the regional layer of the AGF above it, could each hold a defensible claim to a decision — over a shared venture's direction, over which unit's standard applies to a joint activity — and the network's subsidiarity principle, which assigns authority by context and fit rather than by a fixed hierarchy, does not itself say which claim controls when the two are contested rather than merely overlapping. The decision the two centers disagree about does not get made until one center's claim is settled by some means outside the disagreement itself.

Sources: chatgpt (“Jurisdictional Overlaps and Conflicts”); claude (“Potential for conflict and stalemate”); pi (“Overlapping jurisdictions and unclear boundaries”).

9.5.8 Accountability Loss Through Opacity

What it is. Where a decision is the product of several governing centers acting in sequence or in combination, no single center can be held to account for the decision’s outcome, because no single center produced it. Responsibility diffuses across the centers that contributed, and the diffusion is genuine rather than evasive: each center’s individual contribution may be fully defensible, and the outcome still traces to no one center’s decision alone.

Structural source. Polycentric governance produces outcomes through the interaction of many centers rather than through a single chain of command, and that interaction is what makes the system’s outcomes traceable to a combination of decisions rather than to a decision. A system built this way has no natural point at which to locate accountability for a combined outcome unless it adds a mechanism for that purpose specifically, since the structure that produces the outcome does not by itself produce a record of which center’s contribution mattered how much.

In a HAO-shaped network. An outcome shaped by a UME’s local decision, a SEP’s joint governance, and a ruling from a layer of the AGF above both could leave a member harmed by the outcome with no single body to hold accountable for it, since each contributing center can correctly report that its own decision, taken in isolation, was within its mandate. Where responsibility is contested this way, the system’s own account of what happened can shift between the contributing centers faster than any external observer can verify it.

Sources: chatgpt (“Accountability and Transparency”); claude (“Complexity and opacity”); gemini (“Blurred Lines of Responsibility,” “Blame Shifting”).

§9.6 Countermeasures and Detection

This section catalogs the mechanisms the source material proposes against the patterns cataloged in §9.1 to §9.4 and the failure modes enumerated in §9.5. Entries are organized by what a mechanism does rather than by which source names it. §9.6.1 to §9.6.4 treat detection: mechanisms that look for a pattern already running. §9.6.5 to §9.6.10 treat structural resistance: mechanisms that change the arrangement of roles, bodies, or exit rights so that a pattern is harder to sustain. §9.6.11 to §9.6.15 treat practice and culture: mechanisms that act on what members are trained to notice and on what the network records about its own failures. §9.6.16 to §9.6.19 hold four mechanisms added after the first draft of this section: three recovered by the sweep described below, and one the extraction had recorded without the target its source states for it. They are grouped at the end rather than interleaved with the functional groups above. Two of the four are detection mechanisms and belong with §9.6.1 to §9.6.4; one is a set of tracked indicators that no source presents as a countermeasure at all; the fourth is an onboarding measure and belongs with §9.6.11 to §9.6.15. §9.6.20 names the governance instruments the extraction records that belong to chapter 03 rather than to this chapter. §9.6.21 states, for every threat in §9.1.2 and every failure mode in §9.5, what the material proposes against it.

Mechanism names are the sources’ own, on the same basis as the threat names in §9.1: they are retained so that a later reader can trace an entry back to the material that produced it. Five limits apply to everything in this section, and the second governs how §9.6.21 should be read.

1. **Nothing here has been built or tested.** No mechanism below has been implemented in an operating network, and no source reports an outcome from one. Each is a proposal recorded in a red-team exercise conducted against the model’s documentation. The limits stated in §9 apply to this section as they do to the threat catalog, and the direction of the error is not known: an untested countermeasure may fail against the pattern it names, and it may also work against patterns nobody proposed it for.
2. **The threat-to-countermeasure mapping is the extraction’s and this chapter’s, not the sources’.** In the source documents the countermeasures appear as lists attached to a response’s overall subject rather than as a mapping onto named threats. Five attachments are made at the source level, counting an attachment as source-level only where a source itself states a mechanism against a named pattern or scenario. Three of the five are the inline **Prevention:** lines beneath the timeline items in p13r01, Response 1; the fourth is the set of responses the worked scenario document states beneath the case it develops; the fifth is a glossary definition that names its own target. Three of the five land on entries in the catalog, and each is marked where it occurs: the language watchdogs of §9.6.3 against the pattern §9.4.9 names Lexical Drift; the three responses in §9.6.1 and §9.6.10 against the scenario treated in §9.3; and the ethical stress tests noted in §9.6.15 against Mission Creep. The other two, at §9.6.16 and §9.6.19, are attached by their source to scenarios that §9.1.2 does not catalog, so neither supports a row in §9.6.21 by itself; §9.7.7 treats the question they raise. A heading can also frame a whole block of mechanisms around one pattern without stating any of them against it individually — p12r01’s third response does this for Mission Creep — and framing of that kind is not counted here. Every other correspondence recorded in §9.6.21 was drawn in the staged extraction, in the primary-source sweep described below, or in this chapter, and the cells state which. A correspondence marked as inferred is a judgment that a mechanism bears on a threat; it is never a claim that a source says so.
3. **The polycentric countermeasure material is general governance content.** The extraction records that the pass producing the mechanisms in §9.6.20 largely reproduces general public-administration material rather than mechanisms specific to this model, and that four independent model variants converged on similar generic answers to a prompt asking for novel solutions. That observation is the extraction’s. No claim is made here that any named study or research tradition established these instruments: as with the failure modes in §9.5, the sources for them are model-generated responses citing no author or paper.
4. **The mechanisms name no one to run them.** Several entries below specify an overseeing role — a watchdog, a parallel board, an analytical system reviewing decisions — without specifying who appoints it, who funds it, or what holds it to account. One source raises the question for its own proposal and supplies no answer. The gap is recorded here and carried forward to §9.7.
5. **Several are posed as questions rather than specified.** A number of these mechanisms appear in the sources under headings for research directions, and are stated as questions about what could be built or modeled rather than as descriptions of something specified. The entries below mark where that is so. Where a mechanism is posed as a question, the description here reports the question’s subject and does not convert it into a specification the source does not contain.

Provenance of this section’s material. This section was first written from the staged extraction alone, as the rest of chapter 09 was. That extraction was then found to be incomplete against the documents it was drawn from: at least one mechanism carrying a source-level attachment had not been carried into it, and

several mechanisms it did record had been separated from threats their own wording bears on. A targeted sweep of the primary sources was run in consequence, reading both adversarial interview passes, the worked scenario document, all four polycentric-governance responses, and both glossary files in full, and reporting only what the extraction had missed or misattributed. Three results follow, and a reader should be able to tell which claims rest on which body of material.

- Three mechanisms are new to this section and came from the sweep rather than from the extraction: §9.6.16, §9.6.17, and §9.6.18. Their entries say so. A fourth, §9.6.19, was in the extraction as a mechanism with no stated target, and the primary source turns out to state one; that entry says so too.
- Five of the threats the sweep examined, each of which the extraction left with nothing, are recorded here as partial: Divide and Conquer, Manufacturing Dissent, Poisoning the Well, Exploiting Blind Faith in Tech, and Hijacking the Support Network. Trojan Horse stands in the same position — the extraction attaches nothing to it either, and its partial status rests on the correspondence §9.6.2 draws — but it was not among the entries the sweep was asked to check. Three of the five moved status in this section — the other two were already partial on the correspondence §9.6.2 draws, and the sweep added a second mechanism to each. Every one of the five rests on an attachment the sweep marks as inferred rather than stated by a source, and every cell in §9.6.21 says so in those terms. None of the five is a source-level attachment, and partial here records that a mechanism bears on the threat, not that anyone proposed it for the threat.
- Six threats were checked against every primary source and confirmed to have nothing directed at them. Four came out of the first sweep: Founder Capture, Talent & Idea Poaching, Greenwashing Through Imitation, and Inaction as a Weapon. A supplementary sweep covered the two the first had passed over, Tyranny of Small Kindness and Weaponized Burnout, and found the same. Their cells in §9.6.21 state the stronger finding this permits — that the sources offer nothing, rather than that the extraction records nothing. Three citation errors in the extraction were corrected at the same time, and are noted at the entries that carry them.

Everything not covered by those three results rests on the staged extraction, on the same footing as the rest of chapter 09.

9.6.1 Ethical Audits Beyond Financial Matters

The mechanism. Audit coverage is extended past financial anomalies to patterns of conduct that work the framework's gray areas for advantage. The stated object of detection is a pattern rather than an incident, and the source's framing of the change is a reorientation of the network's technical capacity toward harmony rather than toward technical optimization alone.

What it detects. Conduct that leaves no financial irregularity because none occurred. The class of patterns this describes is the one §9.3 treats: acts that are lawful, defensible in the network's own vocabulary, and invisible to an audit that examines whether the figures are correct, because the figures are correct.

What it is directed at. The source names it as one of three possible responses to the framework-turned-against-itself scenario (§9.3.1 to §9.3.5), which it presents as unresolved and does not close with any of the three. It states an explicit limit on the first of the three and states none on this one. The extraction also maps it onto Death by a Thousand Papercuts (§9.3.6), on the basis that both patterns stay below the threshold at which any single instance can be established as a breach; that second correspondence is the extraction's rather than the source's.

Sources: `~/code/papr/icn/security-bad-actors-and-mmm.md`.

9.6.2 Algorithm Auditing

The mechanism. Standing scrutiny of the network's decision-support tools, directed at detecting bias in them and at maintaining the condition that they augment rather than replace human judgment and value-based decision-making.

What it detects. Bias in a tool's output, as distinct from error in it. The entry is one line in a terminology list, and the source states no procedure, no cadence, and no party responsible for conducting the audit.

What it is directed at. The extraction names no threat for this mechanism. The correspondence to the three tool-directed patterns — Trojan Horse (§9.4.5), Poisoning the Well (§9.4.6), and Exploiting Blind Faith in Tech (§9.4.7) — is this chapter's, and it rests on those entries' own stated precondition that no periodic bias audit is conducted. The fit is closest for Trojan Horse, where the defect is in the tool. For the other two it is weaker: a mechanism that scrutinizes a tool is not on its face a mechanism that checks the provenance of the data the tool was trained on or the inputs supplied to it at the point of use. §9.6.21 records all three as partial for that reason.

Sources: `~/code/site-provide-coop/!! Localization.md`, term 8. The staged extraction attributes this term to `~/code/site-provide-coop/Values--.txt`, which is an error. The two files hold the same glossary pass, and the primary-source sweep establishes the difference between them: `!! Localization.md` carries terms 6 to 20 in three headed groups, and `Values--.txt` carries only terms 11 to 20, dropping terms 6 to 10 while retaining the section header that belongs above them. Terms 7, 8, and 9 — mandatory psychoeducation, algorithm auditing, and the sanctioned-practice sense of "Trojan Horse" recorded at §9.1.6 — exist only in `!! Localization.md`. The citation here follows the file the term appears in.

9.6.3 Language Watchdogs and Linguistic Immunology

The mechanism. Two forms appear. In the first, a designated role analyzes patterns in word usage across the network: which terms are falling out of use, and which are acquiring a connotation they did not previously carry. In the second, the analysis is conducted by a system trained to detect shifts in language patterns, examining the rhetorical strategy by which a proposal is framed as well as its content.

What it detects. Change in what the network's shared vocabulary requires, at a rate slow enough that no single usage is remarkable. The first form detects the change by comparison over time; the second is proposed as a research direction rather than as a specified system, and the source states it as a question about what could be trained rather than as a description of something built.

What it is directed at. Lexical Drift (§9.4.9). This is one of the five attachments made at the source level: the source attaches the watchdog role to the pattern it names “Capturing the Narrative”, which is the pattern §9.1.2 catalogs as Lexical Drift. The same prevention line carries two further measures, recorded here so that the attachment is stated in full: mandatory participation in ethics debate clubs, and the celebration of dissenters in the network’s official history. Neither detects a shift in usage, and the status §9.6.21 records for Lexical Drift rests on the watchdog role.

Sources: p13r01, Response 1 (language watchdogs) and Response 3 (linguistic immunology).

9.6.4 Anomalous-Pattern Review

The mechanism. Two proposals treat an anomaly in the network’s own records as the object of review. In the first, analytical systems examine network data for hidden patterns indicating factional alignment, including alignment among participants who intend no harm; the source adds that such systems are themselves to be kept under scrutiny, so that they do not become an instrument of control in their own right. In the second, members hold positions on the outcomes of pending decisions, and unusual movement on a proposal whose outcome appeared settled is treated as a signal that the decision warrants closer examination. The second is posed in its source as a question about whether such a market would expose manipulation, and no procedure for running one is stated.

What it detects. In the first case, coordination that no participant has declared and that formal records do not show. In the second, a divergence between the expected and the priced outcome of a decision, which is a signal that something not visible in the deliberation is affecting it.

What it is directed at. The extraction records both mechanisms without a stated target, and this chapter does not supply one. Both are listed in the extraction among mechanisms it marks as lower priority. The self-scrutiny condition attached to the first is the point at which one source raises the question of who oversees the overseers, and it names no mechanism for that scrutiny; §9.7 carries the question forward.

Sources: p13r01, Response 3 (analytical review of factionalism patterns) and Response 2 (positions held on decision outcomes).

9.6.5 Rotation as a Principle

The mechanism. No member holds the same role for an extended period. The source names leadership positions, resource-allocation committees, and prominent roles in conflict resolution as within scope, and states two effects: entrenchment is prevented, and the range of skills held across the network widens.

What it prevents. Accumulation of influence in a position over time. The mechanism acts on tenure rather than on conduct, so it does not require that anything be detected before it operates.

What it is directed at. The extraction maps it onto Hidden Oligarchy (§9.2.1), Cult of Personality (§9.2.2), and Founder Capture (§9.2.9). The mapping to Founder Capture is where the mechanism’s limit is clearest, and §9.2.9 states it: rotation reaches formal office, and the authority a founder holds is informal, so a rotation requirement can run in full while the influence the pattern converts is untouched by it. This is the first of the three mechanisms §9.2.9 refers to as general anti-entrenchment mechanisms cataloged here.

Sources: p13r01, Response 1.

9.6.6 Shadow Governance and Shadow Boards

The mechanism. Two forms appear. In the first, units selected at random form temporary parallel structures for consequential decisions, and the recommendations those structures produce are compared against the outcomes the official process reached. In the second, every major decision body has a standing unofficial counterpart whose members are selected at random, whose terms are strictly limited, and which holds veto standing over decisions it judges too risky.

What it detects or prevents. The first form detects divergence: a systematic gap between what a parallel body drawn by lot recommends and what the official body decides is a signal about the official body rather than about any single decision. The second form does not detect, it blocks, and the source states its purpose as disrupting the ability of a faction to hold entrenched control.

What it is directed at. The extraction maps both forms onto Hidden Oligarchy (§9.2.1), Weaponized Consensus (§9.2.3), and Founder Capture (§9.2.9). Against Weaponized Consensus the mapping is partial in a specific way: the pattern in §9.2.3 works by extending deliberation without a terminating condition, and neither form of this mechanism states how a parallel body brings a deliberation to an end. This is the second of the three mechanisms §9.2.9 refers to, and it is the one whose selection by lot §9.2.9 describes.

Sources: p13r01, Response 1 (shadow governance) and Response 3 (shadow boards).

9.6.7 Values-Based Circuit Breakers

The mechanism. Metrics tied to long-term wellbeing carry thresholds, and a breach triggers an audit automatically and mandates specified actions, without regard to the unit’s financial performance. The source names two thresholds: disparity in member access to mental-health resources, and leadership composition skewing toward a single demographic group.

What it detects. A condition rather than an act. The mechanism requires no finding about anyone’s conduct and no judgment that a pattern is under way: the threshold is either breached or it is not, and the response follows from the reading.

What it is directed at. The extraction maps it onto Mission Creep (§9.1.3) and, through the leadership-composition threshold, onto Hidden Oligarchy (§9.2.1).

Sources: p12r01, Response 3.

9.6.8 Breakaway Provisions

The mechanism. Units pre-negotiate the terms on which they may exit, against the case that the wider network departs unacceptably from its stated commitments. The source states two effects: the provision deters, and where deterrence fails, the network fragments into clusters that retain the commitments rather than being held whole under the control of the parties that captured it.

What it prevents. It does not detect. It changes what a capture is worth by limiting what a captured network retains, and the change operates whether or not anyone identifies the capture as it occurs.

What it is directed at. The extraction maps it onto Founder Capture (§9.2.9) and onto the escalation of Mission Zealot (§9.1.5, §9.2.8) and Purity Trap (§9.1.4) — that is, onto the case in which those patterns have already carried the network past a point the exiting unit will accept, rather than onto the patterns as they operate. This is the third of the three mechanisms §9.2.9 refers to.

Sources: p12r01, Response 3.

9.6.9 Controlled Burns and Virtuous Viruses

The mechanism. Two proposals introduce change into the network's own arrangements on a schedule the network sets rather than in response to an event. Controlled burns inject low-level disruption deliberately — simulated conflicts, temporary limits on access to a resource, the sudden dissolution of an established unit — each followed by a debrief. Virtuous viruses model elements of the network's operating arrangements on processes that spread, mutate, and recombine, so that no part of the network stays fixed long enough for a faction to hold it; the source poses this one as a question about whether such modeling is possible, and specifies no element to which it would apply.

What it prevents. In both cases, the persistence a capture depends on. Neither mechanism identifies a party or an act; both act on the stability that any pattern requiring accumulated position would need.

What it is directed at. The extraction maps virtuous viruses onto Hidden Oligarchy (§9.2.1), Cult of Personality (§9.2.2), and Founder Capture (§9.2.9). It maps controlled burns onto resilience generally rather than onto a named pattern, and notes that they connect to the antifragility scenario modeling specified in `docs/15-advanced/06-systems-simulation-and-stress-testing.md`, which §9 cites as existing coverage rather than restating.

Sources: p13r01, Response 1 (controlled burns) and Response 2 (virtuous viruses).

9.6.10 Rule Revision and Redefined Member Wellbeing

The mechanism. Two of the three responses the source names to the framework-turned-against-itself scenario. The first is immediate revision of the rule to close the specific gap that was used. The second redefines what the network counts as member wellbeing so that it includes the psychological cost borne by the members of a unit asked to rationalize conduct they judge harmful, which relocates responsibility for that conduct from market conditions to the network's own participants.

What they prevent. The first removes one instance and nothing else. The source states its own limit: it buys time, and where it is not accompanied by wider reform it generates resentment. §9.3.5 states a further limit that follows from the vulnerability it names — a closed gap adds a rule, the rule joins the set being optimized against, and the process continues from the revised set. The second changes a measure rather than a procedure, and the source states it as a question about what the network's wellbeing measures could be made to include rather than as a specification of how.

What they are directed at. The framework-turned-against-itself scenario (§9.3.7), at the source level. The source presents the scenario as unresolved by design and offers none of its three responses as a resolution.

Sources: `~/code/papr/icn/security-bad-actors-and-mmm.md`.

9.6.11 Devil's Advocate Simulation and Mandatory Apprenticeship

The mechanism. Two forms. In the first, a rotating group is assembled on the model of a corporate red team and tasked with devising strategies for turning a live unit proposal to harm, with the exercise directed at morally ambiguous conduct rather than at conduct that is plainly unlawful; the wider network then has to identify the vulnerability and answer it. In the second, a period of service in a designated contrarian role is a precondition of holding a leadership position. The source states the second form's effect on the leadership pipeline directly: it removes the standing of a leadership group defined by moral purity.

What it prevents. Exposure to the ambiguous case is the stated object in both forms. The first produces a finding about a specific proposal; the second produces a population of leaders who have had to occupy the adversary's position before holding office.

What it is directed at. The extraction maps the pair onto Mission Creep (§9.1.3), Purity Trap (§9.1.4) through the forced engagement with ambiguity, and Founder Capture (§9.2.9) through the leadership-pipeline requirement. The Founder Capture mapping is subject to the limit §9.2.9 states, since a service requirement conditions formal office and not informal standing.

Sources: p12r01, Response 3 (devil's advocate simulation); p13r01, Response 3 (mandatory devil's advocate apprenticeships).

9.6.12 Branching Corruption Exercises

The mechanism. Members work through branching scenarios in which ethical compromises are presented as growth opportunities. Choices carry simulated consequences for network health, reputation, member wellbeing, and financial return. The exercise also records the justification a participant gives for each

choice, and the source treats that record, rather than the choice, as the output: it allows common rationalization patterns to be identified across participants and answered through targeted education before they are used in a live decision.

What it detects. Rationalization patterns in the population, as distinct from a disposition in any one participant. The unit of analysis is the reasoning offered, and it is collected under conditions where offering it costs nothing.

What it is directed at. The extraction maps it onto the rationalization patterns of Mission Creep (§9.1.3) generally rather than onto a specific sub-pattern of it. In the source it appears among proposals directed at Mission Creep as a whole.

Sources: p12r01, Response 3.

9.6.13 Contests of “Elegant Evil”

The mechanism. A gamified challenge invites members to design contract gaps that violate the network’s values while remaining likely to survive legal scrutiny, or to devise ways of working a regulatory process. The stated purpose is to have the network find its own pathways to corruption before an outside party does.

What it detects. Vulnerabilities in the network’s own drafting and in its regulatory position, found by directed search rather than by monitoring. The mechanism produces a list of available moves and nothing about whether anyone has made one.

What it is directed at. The extraction maps it onto the framework-turned-against-itself scenario (§9.3), where the object of the exercise and the mechanism of the threat are the same — a gap that is lawful to use — and onto Regulatory Capture (§9.4.1) through the part of the challenge concerning regulatory process. The second mapping is partial in a way worth naming: Regulatory Capture as §9.4.1 states it is conducted by incumbents from outside the network, and an internal contest that maps the network’s regulatory exposure produces knowledge of that exposure rather than a defense against the party exploiting it.

Sources: p12r01, Response 3.

9.6.14 Myth of the Fallen

The mechanism. Where a unit or venture has been captured, the case is written up honestly and shared openly within the network as an educational document. The stated purpose is a standing record that good intentions confer no immunity, together with a catalog of the early indicators the case exhibited.

What it prevents. Nothing directly. Both sources that name it treat it as a cultural mechanism directed at complacency rather than as a detection mechanism, and the extraction records it that way.

What it is directed at. No threat by name. The extraction records it as general.

Sources: p12r01, Response 3; and independently as term 19 in `~/code/site-provide-coop/!! Localization.md` and `~/code/site-provide-coop/Values--.txt`. Term 19 falls in the range both files carry, and the sweep confirms both citations.

9.6.15 Further Mechanisms Named Without a Stated Target

The extraction records a further set of mechanisms that no source attaches to a named threat, and marks them as lower priority for this chapter. One entry in the list below is an exception and is marked as such, and one that the extraction placed here has been moved out to §9.6.19, its source having stated a target for it. The rest are listed here so that the coverage in §9.6.21 can be read as a statement about the mapping rather than about the inventory: a threat recorded there as having no countermeasure proposed may still fall within the reach of something in this list, and no source says so. Several of the entries below appear in their sources under headings for research directions and are posed as questions, per limit 5 above; the descriptions state what each question is about and do not supply a specification.

- A network immune system: self-correcting mechanisms triggered at thresholds of stress, conflict, or detected deviation, combining human-led intervention with algorithmic detection. The extraction records an unresolved question about this term, since one source attributes the function to a specific entity — term 11 of the same glossary lists an “immune system” for the network among provide.io’s functions — and another poses it as an open research direction; §9.7.4 carries the question forward.
- Mandatory psychoeducation at onboarding, in conflict resolution and communication, stated as preventive rather than remedial. This is term 7, and it exists only in `!! Localization.md`; the extraction’s citation of `Values--.txt` for it is the third of the three errors §9.6.2 records.
- Ethical stress tests: simulations exposing vulnerabilities in the model. This entry is one of the five source-level attachments noted in the opening limits, and the only one in this list — the source states their aim as uncovering the potential for gradual erosion of the mission rather than overt ethical failure, which attaches them to Mission Creep (§9.1.3).
- Mandated outsiders: a rotating skeptic drawn from another unit, joining temporarily with the task of identifying potential for conduct contrary to the network’s commitments.
- Red teams assembled before any major initiative, tasked with finding how the plan could be turned to advantage, as a generalization of the mechanism in §9.6.11.
- Real-time dilemma simulation, in which members respond to escalating moral dilemmas without preparation and the responses are examined by the group afterward.
- Public post-mortems on close calls, in which an exposed attempt is written up in anonymized form as an educational document, on the model of §9.6.14 but covering attempts rather than completed captures.

- An adversarial analytical system maintained separately from the network’s collaborative tools and tasked with finding vulnerabilities and running simulated attacks. The source names the attack classes it would simulate, and one of them is a targeted misinformation campaign, which is the mechanism Divide and Conquer (§9.4.4) works by. The extraction’s paraphrase dropped that detail; the sweep restored it, and §9.6.21 records the resulting attachment as inferred.
- Game-theoretic modeling directed at making conduct aligned with the network’s stated values the rational choice. This overlaps the game-theoretic material already specified in `docs/15-advanced/06-systems-simulation-and-stress-testing.md`, and §9 records that no source connects that material to any threat in this catalog.

Sources: p12r01, Responses 1 to 3; p13r01, Responses 1 to 3; `~/code/site-provide-coop/!! Localization.md` (terms 7, 11, 16, 17); `~/code/site-provide-coop/Values---.txt` (terms 11, 16, 17 only, per §9.6.2).

9.6.16 Transparency Rotation, Public Decision Dashboards, and Gamified Detection

The mechanism. A single prevention line in the source combines three measures: a mandatory transparency rotation on the committees that allocate resources; public dashboards on which the justification given for a decision is recorded and can be tracked; and gamified challenges in which members analyze the network’s own data to identify resource manipulation presented as legitimate oversight.

What it detects. Manipulation of resource allocation conducted in the form of oversight. The three measures act at different points: the rotation changes who sits on the committee, the dashboard makes the stated justification for a decision available for comparison against the decision, and the challenge distributes the analysis across the membership rather than concentrating it in a review body.

What it is directed at. The source attaches this line to a scenario of its own — a coalition circulating accounts of bias in resource-allocation decisions, algorithmic and human, in order to degrade trust in the network’s leadership. The attachment is source-level, but the scenario it attaches to has no entry in the catalog at §9.1.2; it is one of the two source-level attachments in that position, the other being §9.6.19. The sweep recorded in the provenance note above finds that scenario thematically close to Manufacturing Dissent (§9.4.8) without being identical to it: Manufacturing Dissent works by manipulating what the health-monitoring system displays, and this scenario works by circulating accounts about decisions. The link to Manufacturing Dissent is therefore inferred rather than stated, and §9.6.21 marks it as such.

Sources: p13r01, Response 1. Recorded in the primary-source sweep; absent from the staged extraction.

9.6.17 Early-Warning Flagging of Unusual Activity

The mechanism. Data on communication patterns, resource accumulation, and the justifications recorded for decisions is analyzed to highlight unusual clusters of activity for human examination. The source states a limit as part of the mechanism: the analysis cannot establish intent, and what it produces is a starting point for an investigation rather than a substitute for one.

What it detects. A cluster rather than a conduct. The output identifies where to look, and the source assigns the question of whether anything is there to a person.

What it is directed at. The source lists it in a general prevention block alongside rotation, shadow governance, and controlled burns, and attaches it to no named threat. Its stated limit — that it does not replace human review — bears on the mechanism of Exploiting Blind Faith in Tech (§9.4.7), which is a system’s output being acted on without independent checking. That attachment is inferred by this chapter and by the sweep that supplied the entry; the source names no threat next to the mechanism.

Sources: p13r01, Response 1. Recorded in the primary-source sweep; absent from the staged extraction.

9.6.18 Ethical Health Metrics

The mechanism. A set of tracked indicators of the network’s condition extending past financial data, divided into quantitative measures of resource-distribution equity and qualitative measures drawn from member surveys on belonging and fairness.

What it detects. Nothing on its own. The entry is a definition in a terminology list, and the source does not present it as a countermeasure or state any use to which the indicators would be put.

What it is directed at. No threat. The sweep that supplied this entry rates it the least confident item it found and marks the link doubly inferred: from a glossary term to a countermeasure function, and from that function to a threat. The reasoning it records is that independently tracked measures of resource distribution are the kind of record against which a fabricated distribution claim could be checked, which would bear on Manufacturing Dissent (§9.4.8). That rating is reproduced here as given and is not strengthened, and §9.6.21 does not rest Manufacturing Dissent’s status on this entry alone.

Sources: `~/code/site-provide-coop/!! Localization.md`, term 6. Recorded in the primary-source sweep; absent from the staged extraction.

9.6.19 Buddy System, Values Quizzes, and Fallacy Training at Onboarding

The mechanism. A single prevention line in the source combines three measures applied at the point where a unit is formed: mentors for newly formed units drawn from a pool by random assignment; values assessments during onboarding, set as case studies; and early training in identifying the logical fallacies used to justify conduct at odds with the network’s commitments.

What it prevents. Misrepresentation of the network’s arrangements to a unit that has no independent basis for checking the account it is given. The three measures act at different points: random assignment removes a mentor’s ability to select whom they mentor, the case-study assessment gives a new unit a reference against which a mentor’s account can be compared, and the fallacy training acts on the form of a justification rather than on its content.

What it is directed at. A scenario the source states directly above the prevention line: a member with facility in rhetoric presenting as a mentor to newly formed units, misrepresenting the network’s policies and directing those units toward personal connections rather than official resources. The attachment is source-level. The scenario has no entry in the catalog at §9.1.2, so this mechanism supports no row in §9.6.21 by itself, and it is the second of the two source-level attachments in the material that land on nothing the catalog names; §9.7.7 treats the question that raises. No correspondence to a cataloged entry is drawn here.

Sources: p13r01, Response 1. The staged extraction recorded the buddy system among the mechanisms it marks as having no stated target, and dropped the values quizzes, the fallacy training, and the scenario the source attaches all three to; the attachment is restored here from the primary source.

9.6.20 Governance Instruments That Belong to Chapter 03

The extraction records a further set of mechanisms proposed against the polycentric failure modes of §9.5. They are governance design rather than detection or structural resistance: they specify how coordinating bodies are constituted, what agreements between units contain, and how authority over a contested decision is settled. Chapter 03 specifies affirmative governance design for a Humanized Autonomous Organization (HAO) — the network’s coordinating framework — and this chapter cross-references it rather than restating it, per the boundary §9 sets. The instruments are named below with the failure modes the extraction maps them to, and are not specified here.

- Bridging organizations, mapped to the coordination cluster, which §9.5 divides between Absence of Overarching Vision (§9.5.1) and Cross-Boundary Problem Paralysis (§9.5.2).
- Joint governance agreements, mapped to the jurisdictional-overlap cluster, which §9.5 treats as Conflict Deadlock (§9.5.7), and bearing on Cross-Boundary Problem Paralysis (§9.5.2) through the cross-boundary responsibilities they record.
- Equity scorecards paired with redistributive transfers, mapped to the equity and representation cluster, which §9.5 divides between Elite Capture of Governing Units (§9.5.4) and Equity Drift (§9.5.5).
- Citizen audits with open-data tracking of decisions and outcomes, mapped to the accountability and transparency cluster, which §9.5 treats as Accountability Loss Through Opacity (§9.5.8).
- Policy labs and governance sandboxes with sunset clauses, mapped to the entrenchment and path-dependence cluster, which §9.5 folds into Absence of Overarching Vision (§9.5.1).
- Resilience networks, mapped to resilience across the modes rather than to one of them, and overlapping the containment protocols in `docs/06-lifecycle/03-collapse-and-containment-protocols.md` and the antifragility modeling in `docs/15-advanced/06-systems-simulation-and-stress-testing.md`.
- Social impact bonds for governance, recorded by the extraction as a funding mechanism directed at no particular failure mode.

The staged extraction omitted the fourth polycentric response, `005-r-pi.md`, which states six further instruments. That response differs from the other three in form: each instrument is stated against a named gap in the construction “To address X...”, and the gap terms it names are the ones the same variant supplied to the earlier gaps prompt. The correspondence from an instrument to a gap term is therefore the source’s; the correspondence from that gap term to a failure mode is §9.5’s consolidation, which reads that variant’s terms at §9.5.1, §9.5.2, §9.5.5, §9.5.6, and §9.5.7.

- Knowledge-sharing platforms for data and practice among centers, stated against information asymmetry, and adaptive and participatory monitoring and evaluation, stated against the problem of assessing a polycentric system’s performance. §9.5.6 reads both of those terms for Monitoring and Enforcement Cost at Scale.
- Capacity-building initiatives — training, technical assistance, or financial support directed at less-resourced centers so that they can take part in decisions — stated against capacity and resource constraints, which §9.5 reads for its resource-provision sense at Absence of Overarching Vision (§9.5.1) and for its participation sense at Equity Drift (§9.5.5).
- Mediation and conflict-resolution mechanisms, stated against conflicting interests and values, and the clarification of each center’s roles, responsibilities, and decision-making authority, stated against overlapping jurisdictions and unclear boundaries. §9.5 treats the second of those terms as Conflict Deadlock (§9.5.7).
- Multi-level coordination bodies, stated against scaling up and coordination across scales, which §9.5 treats as Cross-Boundary Problem Paralysis (§9.5.2).

Two of the eight failure modes receive nothing directed at them in this set: Race to the Bottom Between Jurisdictions (§9.5.3), and — except through the resource-allocation material folded into it — the system-wide provision aspect of Absence of Overarching Vision (§9.5.1).

Limit 3 in the opening applies to this subsection in particular. The extraction’s own assessment is that the material it carried largely reproduces general public-administration content rather than mechanisms specific to this model, and that it is more useful for its taxonomy of failure modes than for its mechanism language. The fourth response, recovered here, is of the same kind. §9.5 uses the material for the taxonomy; this subsection records the mechanisms without adopting them.

Sources: `~/code/papr/polycentric-governance/005-r-claude.md`, `005-r-gemini.md`, `005-r-chatgpt.md`, `005-r-pi.md`. The last is absent from the staged extraction and is read here from the primary source.

9.6.21 Coverage

The tables below state, for every entry in the catalog at §9.1.2 and every failure mode in §9.5, what the material proposes against it. Status carries three values:

- `addressed` — one or more mechanisms in this section are directed at the entry in the material.
- `partial` — a mechanism bears on the entry, but the material directs it at a wider or adjacent target, or the source does not offer its own answer as a resolution. Cells reaching this status on an attachment nobody stated say so in the cell, in the form “attachment inferred”.

- **none proposed** — nothing in the material is directed at the entry. Where the primary-source sweep checked the entry directly, the cell says so, and the finding is then about the sources rather than about the extraction.

Two properties of the tables should be read together with limit 2 above. First, a status is a statement about the mapping recorded in the material, not about whether a mechanism would work: nothing here has been tested, and **addressed** records that something was proposed, not that the pattern is answered. Second, every correspondence this chapter or the sweep draws rather than inherits from a source is marked as inferred in the cell where it appears, and the count of them exceeds the count of source-level attachments several times over.

The failure modes are tabulated separately from the threats. §9.1.1 defines the polycentric category as the one requiring no actor, so the column head “Threat” does not apply to them.

Threat	Countermeasure(s)	Status
Mission Creep	Ethical stress tests (§9.6.15), whose glossary entry names this pattern as their target; devil’s advocate simulation and mandatory apprenticeship (§9.6.11); branching corruption exercises (§9.6.12); values-based circuit breakers (§9.6.7). The status rests on the first, which carries the source-level attachment; the other three are the extraction’s correspondences	addressed
Purity Trap	Devil’s advocate simulation and mandatory apprenticeship (§9.6.11), through forced engagement with ambiguity; breakaway provisions (§9.6.8), mapped to the escalation of the pattern rather than to the pattern. Both correspondences drawn by the extraction from prevention blocks their source attaches to no named pattern	partial
Mission Zealot	Breakaway provisions (§9.6.8), mapped to the escalation of the pattern rather than to the pattern	partial
Hidden Oligarchy	Rotation as a principle (§9.6.5); shadow governance and shadow boards (§9.6.6); values-based circuit breakers (§9.6.7); virtuous viruses (§9.6.9). All four correspondences drawn by the extraction from prevention blocks and research prompts their sources attach to no named pattern	partial
Cult of Personality	Rotation as a principle (§9.6.5); virtuous viruses (§9.6.9). Both correspondences drawn by the extraction from material its source attaches to no named pattern	partial
Weaponized Consensus	Shadow governance and shadow boards (§9.6.6), which state no terminating condition for a deliberation	partial
Tyranny of Small Kindness	None. A supplementary sweep checked every primary source; the pattern is named once, with nothing attached to it, and the nearest mechanism is stated generally enough to bear on any threat	none proposed
Founder Capture	None directed at the pattern, confirmed against every primary source. The extraction maps five mechanisms onto it — rotation as a principle (§9.6.5), shadow governance and shadow boards (§9.6.6), breakaway provisions (§9.6.8), virtuous viruses (§9.6.9), and the devil’s advocate pair (§9.6.11) — and each reaches formal office, a static target, or the leadership pipeline, rather than the informal standing the pattern converts (§9.2.9). The sweep checked both interview passes, the worked scenario, all four polycentric responses, and both glossary files for a mechanism directed at a founder’s informal legitimacy and found none; the worked scenario poses the case at its closing line and stops without invoking any element of the framework to answer it (§9.3.7, §9.2.9)	none proposed
Weaponized Burnout	None. A supplementary sweep checked every primary source; the nearest mechanism detects factionalism rather than how work and emotional labor are allotted, which is what the pattern’s precondition names	none proposed
Hijacking the Support Network	Anomalous-pattern review (§9.6.4), whose stated object is factional alignment, of which a personal loyalty network built through manufactured crises is one form; public post-mortems on close calls (§9.6.15). Attachments inferred by the sweep and this chapter; the source’s prevention block covers three threats without tagging any of them, and it never names this one next to either mechanism	partial
Inaction as a Weapon	None in any primary source, confirmed by the sweep: nothing in the material targets covert passivity as a pattern in its own right, and the nearest candidates are an open research question and a mechanism aimed at a single initiative’s exploitability. §9.2.4 records separately that its first precondition is the absence of a mechanism docs/10-evaluation/02-participation-quality-and-alignment-audits.md specifies	none proposed
Framework-turned-against-itself (loophole exploitation)	Ethical audits beyond financial matters (§9.6.1); rule revision and redefined member wellbeing (§9.6.10); contests of “elegant evil” (§9.6.13). The source presents the scenario as unresolved by design and offers none of its three responses as a resolution (§9.3.7)	partial
Death by a Thousand Papercuts	Ethical audits beyond financial matters (§9.6.1), mapped by the extraction rather than by the source	partial
Regulatory Capture	Contests of “elegant evil” (§9.6.13), which map the network’s exposure without answering the party exploiting it	partial
Greenwashing Through Imitation	None in any primary source, confirmed by the sweep: nothing in the material addresses distinguishing the network from firms that adopt its language without its structures. The one glossary term touching public positioning concerns the network’s own policy advocacy	none proposed
Talent & Idea Poaching	None in any primary source, confirmed by the sweep: nothing in the material addresses retaining members against a competitor’s financial incentives, and the research avenues sitting beside this pattern in its own source address data governance and internal groupthink instead	none proposed
Divide and Conquer	The adversarial analytical system (§9.6.15), whose stated remit includes simulating targeted misinformation campaigns, which is this pattern’s mechanism. Attachment inferred by the sweep and this chapter; the mechanism and the pattern appear in the same source response, and the source states the mechanism without naming this threat	partial
Trojan Horse	Algorithm auditing (§9.6.2); the correspondence is this chapter’s, resting on the entry’s own precondition of no periodic bias audit	partial
Poisoning the Well	Algorithm auditing (§9.6.2), a direct-topic match to biased output produced by altered training data; ethical audits beyond financial matters (§9.6.1), extended here from the two threats the extraction attaches it to. Both attachments inferred, by this chapter and by the sweep; neither source names this threat, and algorithm auditing as stated audits the tool rather than its training data	partial
Exploiting Blind Faith in Tech	Algorithm auditing (§9.6.2); early-warning flagging of unusual activity (§9.6.17), whose stated limit is that it starts an investigation rather than replacing one, which is the condition this pattern needs absent. Both attachments inferred, by this chapter and by the sweep; the sources name no threat next to either mechanism, and algorithm auditing as stated audits the tool rather than the inputs supplied to it	partial
Manufacturing Dissent	Transparency rotation, public decision dashboards, and gamified detection (§9.6.16); ethical health metrics (§9.6.18), which the sweep rates its least confident item and which this status does not rest on. Attachments inferred; the source attaches §9.6.16 to a scenario of rumor-spreading about decisions that §9.1.2 does not catalog, and never names this threat, which is a different source’s coinage for manipulation of what the monitoring system displays	partial
Lexical Drift	Language watchdogs and linguistic immunology (§9.6.3), attached to the pattern at the source level, together with the two further measures on the same prevention line	addressed
Failure mode	Countermeasure(s)	Status
Absence of Overarching Vision	Bridging organizations; policy labs and governance sandboxes with sunset clauses (§9.6.20; chapter 03 specifies both)	partial
Cross-Boundary Problem Paralysis	Bridging organizations; joint governance agreements (§9.6.20; chapter 03 specifies both)	partial
Race to the Bottom Between Jurisdictions	None in the extraction, and none among the six instruments in the fourth polycentric response recovered at §9.6.20	none proposed

Failure mode	Countermeasure(s)	Status
Elite Capture of Governing Units	Equity scorecards with redistributive transfers (§9.6.20; chapter 03 specifies this instrument), mapped to the equity cluster rather than to this mode	partial
Equity Drift	Equity scorecards with redistributive transfers (§9.6.20; chapter 03 specifies this instrument)	partial
Monitoring and Enforcement Cost at Scale	Knowledge-sharing platforms; adaptive and participatory monitoring and evaluation (§9.6.20; chapter 03 does not specify these, see §3.4). Each is stated by its source against a gap term — information asymmetry, and assessing a polycentric system’s performance — that §9.5.6 reads for this mode; the step from gap term to mode is §9.5’s	partial
Conflict Deadlock	Joint governance agreements (§9.6.20; chapter 03 specifies them)	partial
Accountability Loss Through Opacity	Citizen audits with open-data tracking (§9.6.20; chapter 03 specifies this instrument)	partial

What the tables show. Of the 22 cataloged threats, two are recorded as addressed, fourteen as partial, and six as having no countermeasure proposed. Of the eight failure modes, none is recorded as addressed, seven as partial, and one as having none proposed; the seven partial entries all rest on the material §9.6.20 assigns to chapter 03, which is general governance content rather than mechanisms specific to this model.

Fourteen partial entries is a larger figure than the material’s own directedness supports, and the cells state why. Exactly one of the fourteen rests on an attachment a source makes itself: the framework-capture scenario of §9.3, which the source presents as unresolved and answers with three responses it does not offer as a resolution. Six rest on correspondences this chapter or the primary-source sweep drew rather than found — five marked inferred by the sweep, three drawn by §9.6.2, with two rows carrying both. The remaining seven rest on correspondences the staged extraction drew. Three of those seven — Purity Trap, Hidden Oligarchy, and Cult of Personality — rest on the same kind of untagged prevention block that yields partial elsewhere in the table, and are recorded as partial on that basis rather than as addressed. A reader counting what the sources actually direct at a named threat should count the two addressed entries and that single partial, and read the rest as a record of what bears on a threat rather than of what anyone proposed for one.

All six threats with nothing proposed against them were checked directly against every primary source and are empty there. Four came out of the first sweep: Greenwashing Through Imitation and Talent & Idea Poaching, both of which act from outside the network’s boundary against a countermeasure set that is almost entirely internal; Inaction as a Weapon, for which the nearest material is an open research question rather than a proposed mechanism; and Founder Capture. A supplementary sweep covered the remaining two, Tyranny of Small Kindness and Weaponized Burnout, which act on the network’s mutual-support and participation norms. Each is named once in the sources, in a single sentence, with nothing attached to it, and the nearest candidate mechanisms are directed elsewhere: the adversarial-analysis tooling nearest the first is stated generally enough to bear on any threat, and the faction-detection tooling nearest the second acts on factionalism rather than on how work and emotional labor are allotted, which is what the pattern’s own precondition names. §9.7 may state all six at the same strength.

Founder Capture is the finding this section ends on, and the sweep raises it from a property of the extraction to a property of the sources. Every source group was checked for a mechanism directed at the informal legitimacy a founder holds by having founded the network, and none holds one. The worked scenario document reaches the case in its closing line, states that it would require holding members in leadership positions to a higher standard than the one applied generally, and stops without naming a mechanism or invoking any element of the framework. The sources raise the problem and decline to resolve it. That is the whole of the claim, and §9.7 carries it forward in that form.

§9.7 inverts these tables and treats the entries with nothing proposed against them, together with the questions this section leaves open: who conducts the mechanisms in §9.6.1 to §9.6.4 and §9.6.17, which entity performs the network immune-system function named in §9.6.15, and whether the scenarios §9.6.16 and §9.6.19 are attached to warrant catalog entries of their own, since they are the two source-level attachments in the material that land on nothing §9.1.2 names.

§9.7 Open Problems

This section inverts the tables in §9.6.21. Where that subsection states what the material proposes against each entry in the catalog, this one treats the entries the material proposes nothing against, together with the questions §9.6 raises and does not settle. Nothing recorded here is a proposal. Where the text goes beyond what a source states, the elaboration is marked as this chapter’s.

Two items below are properties of the material as a whole rather than gaps against a particular threat. The first is that no element of the framework resolves the framework-capture case, which the source developing that case states in its own terms and presents as a condition of the framework rather than as a defect in it (§9.7.3). The three vulnerabilities that scenario is used to identify — metrics-versus-values divergence (§9.3.3), resilience as private capture (§9.3.4), and optimization pressure (§9.3.5) — are open in the same sense and are treated there. The second is that the coverage recorded in §9.6.21 registers what bears on a threat far more often than what anyone proposed for one (§9.7.8).

Two further observations already recorded in §9.6 are not developed again here. The polycentric countermeasure material reproduces general public-administration content rather than mechanisms specific to this model, per limit 3 of that section and §9.6.20. And no source connects the game-theoretic modeling specified in `docs/15-advanced/06-systems-simulation-and-stress-testing.md` to any threat in this catalog; §9.6.15 records that, and drawing the connection would be new work rather than extraction.

9.7.1 Founder and Insider Capture

Founder Capture is the chapter’s central open problem. Two independent source families reach it from different directions, and neither answers it.

The first is the pair of adversarial interview passes. p13r01 develops the pattern narratively under two names, “Idealist Turned Cynic” and, in the more developed variant, “Disillusioned Savior”, and §9.2.9 catalogs both as one pattern: a founding or long-committed member converts standing informal authority into control of the leadership pipeline and reframes dissent as disloyalty. The second is the worked scenario document. `~/code/papr/icn/security-bad-actors-and-mmm.md` closes by raising the case in which the unit exploiting the legal gap is itself a founding entity, influential within the network’s informal structures, states

that this would require holding members in leadership positions to a higher standard than the one applied generally, and stops. It names no mechanism and invokes no element of the framework to answer it (§9.3.7). One source reaches the problem through an individual's turn against the network; the other reaches it through a unit justifying its own conduct in the network's vocabulary. The two are not cross-referenced anywhere in the material.

The primary-source sweep recorded at §9.6 checked both interview passes, the worked scenario, all four polycentric-governance responses, and both glossary files for a mechanism directed at the informal legitimacy a founder holds by having founded the network. None holds one. The finding is about the sources rather than about the staged extraction.

Why the mechanisms in §9.6 do not reach it. §9.6.21 lists the five mechanisms the extraction maps onto this pattern, and each fails against it for a stated reason. Rotation as a principle (§9.6.5) acts on tenure in a formal role, and the authority the pattern converts is not held in a formal role, so a rotation requirement can run in full while that authority is untouched. Mandatory devil's advocate apprenticeship (§9.6.11) conditions entry to formal office on prior service in a contrarian role, and conditions the same surface. Shadow governance and shadow boards (§9.6.6) supply a parallel body selected by lot whose recommendations are compared against the official body's decisions, or which holds veto standing over decisions it judges too risky. Virtuous viruses (§9.6.9) act on the persistence of any position held long enough to be captured. Breakaway provisions (§9.6.8) do not act on the pattern at all; they change what a capture retains by letting units exit on pre-negotiated terms.

The following reading is this chapter's rather than any source's. Each of those mechanisms terminates in the same place. A rotation requirement has to be applied, a parallel body has to exercise its standing, a veto has to be cast, an exit has to be declared. Every one of those is an act by some body against the people who defined the rules that body applies, and the property the pattern turns on is precisely the informal weight those people carry in the bodies that would have to act. The detection mechanisms of §9.6.1 to §9.6.4 and §9.6.17 do not close the gap either, on the same reading: their output is a finding, a divergence, or a cluster warranting examination, and a finding is not an action. The route from a finding to a consequence runs through a decision body, which is the surface in question.

The material states neither who would hold a leadership member to the higher standard the worked scenario names nor on what authority. Chapter 03 specifies affirmative governance design for a Humanized Autonomous Organization (HAO) — the network's coordinating framework — and nothing in the extraction connects that design to this case. This chapter does not supply the connection, because no source supports one.

9.7.2 The Threats Nothing Is Directed At

Six of the 22 cataloged threats are recorded in §9.6.21 as having nothing proposed against them: Founder Capture, Talent & Idea Poaching, Greenwashing Through Imitation, Inaction as a Weapon, Tyranny of Small Kindness, and Weaponized Burnout. All six were checked directly against every primary source, four in the first sweep and two in a supplementary sweep covering the pair the first passed over. All six therefore stand at the same evidential strength: the finding is that the sources hold nothing directed at them, not that the staged extraction failed to record something.

They also share a property of how they appear. Each is named once, in a single sentence, inside a block that its source does not tag with a countermeasure. The material carries five attachments made at the source level, in which a source states a mechanism against a named pattern or scenario (§9.6 limit 2), and none of the five lands on any of these six. Every countermeasure block that follows them covers several threats at once without distinguishing between them, which is the condition under which §9.6.21 records a correspondence as inferred rather than stated.

Grouped by what the absence has in common, they fall into three sets.

Patterns acting from outside the network's boundary. Greenwashing Through Imitation and Talent & Idea Poaching are conducted by parties the network does not govern, against a countermeasure inventory that is almost entirely internal. Nothing in the material addresses distinguishing the network from firms that adopt its language without its structures, and nothing addresses retaining governance-skilled members against a competitor's financial incentives. The research avenues sitting beside each of them in their own source address other subjects: stress testing, game theory, and adversarial tooling beside Greenwashing Through Imitation, and data governance and internal groupthink beside Talent & Idea Poaching.

Patterns acting on support and participation norms. Tyranny of Small Kindness, Weaponized Burnout, and Inaction as a Weapon each turn on how work, support, or participation is distributed among members rather than on a decision procedure, a tool, or the network's vocabulary. The nearest candidate mechanisms are directed elsewhere: at factional pattern detection, at review of a particular risky decision, at entry to leadership, or at a single initiative's exploitability. For Inaction as a Weapon the nearest material is an open research question rather than a proposed mechanism. §9.2.4 records separately that this pattern's first precondition is the absence of a mechanism that `docs/10-evaluation/02-participation-quality-and-alignment-audits.md` specifies, which is existing coverage rather than a countermeasure in this material.

The informal-legitimacy case. Founder Capture stands alone and is treated in §9.7.1.

Read across the three sets, the following is this chapter's observation rather than a claim any source makes. The mechanisms the material does propose act on four surfaces: tenure in formal office, the conduct of a decision procedure, the design and output of a tool, and the usage of shared terms. None of the six threats above acts on any of those four. They act on the network's external boundary, on the distribution of work and support inside it, and on standing held without office.

9.7.3 The Framework-Capture Case Has No Resolution in the Framework

The source that develops the framework-capture scenario states that no single element of the framework resolves it and that no clear-cut correct ethical position is available within it. The scenario is unresolved by the source's own design: its stated purpose is to force a debate and a systemic reassessment, which the source treats as conditions of the network's long-term health. §9.3.7 records this, and the open ending is preserved here rather than closed.

The structural reason the source gives is that the unit's justification is assembled from three commitments the network itself holds — an obligation to its members' economic security, compliance with the law as the operative standard, and efficient use of an available resource (§9.3.2). Answering it requires ranking

those commitments against one another, and the source states that the framework contains no such ranking. What is open is therefore not a missing rule but a missing ordering among rules the network already has.

The three vulnerabilities the scenario is used to identify are open in the same way, and none of the three has a mechanism directed at it. Metrics-versus-values divergence (§9.3.3) is the case in which measured performance and stated values move in opposite directions while every measurement remains truthful; an audit of the figures returns nothing because the figures are correct, and the network holds no measure of the comparison that would show the divergence. Resilience as private capture (§9.3.4) is the case in which adaptive capacity built to absorb network-level shocks is exercised for one unit's benefit; the distinction is in who benefits, and the mechanism does not record that. Optimization pressure (§9.3.5) is the case in which a capable governance-aware process finds moves that are rule-legal and value-harmful as a matter of course, which sets a limit on rule revision as a response: closing a gap adds a rule, and the rule joins the set being optimized against.

Three responses are recorded against the scenario, and the source offers none of them as a resolution: rule revision closing the specific gap, on which the source states an explicit limit; audits extended beyond financial matters, on which it states none; and a redefinition of member wellbeing that would count the psychological cost borne by the members of a unit asked to rationalize conduct they judge harmful, which the source poses as a question rather than as a specification (§9.6.1, §9.6.10). §9.7.8 records that this is the only entry in the coverage table whose partial status rests on an attachment a source makes itself.

9.7.4 The Network Immune System

The material uses one term for two things and does not reconcile them. This chapter records both and takes no position, because the two usages sit in different source families and neither refers to the other.

In the glossary pass, the function is treated as a capability of a named entity. Term 11 defines provide.io by its multiple functions and lists an “immune system” for the network among them, alongside incubator and resource optimizer, in the present tense and without qualification (`~/code/site-provide-coop/!! Localization.md` , term 11; the same term is also present in `~/code/site-provide-coop/Values--.txt` , which carries terms 11 to 20). Separately, term 17 of the same glossary defines the network “immune system” as the self-correcting mechanisms triggered by certain thresholds, and describes an interplay between human-led intervention and algorithmic detection of negative patterns. Term 17 states what the function consists of; term 11 attributes it to provide.io.

In the adversarial interview material, the same function is an unbuilt research question. p12r01, Response 1 lists “Immune Systems for Networks” among its research avenues and poses it as a question about whether principles drawn from biological systems could be applied to the network's growth strategy in order to develop self-correcting mechanisms triggered at thresholds of stress, conflict, or detected deviation from the network's mission. The passage asks whether such mechanisms could be developed. It does not describe mechanisms that exist or assign them to any party.

§9.6.15 catalogs the mechanism on the terms of the glossary definition and marks the question as unresolved. The difference is not a difference in what the mechanism would do; both descriptions name self-correcting responses triggered at thresholds, combining human intervention with algorithmic detection. The difference is in status and in ownership: whether the function is one an existing entity already performs, or a research direction nobody has built. Which reading is correct determines whether the mechanism can be cited as available coverage against any threat in §9.1.2, and no source in the extraction has standing over both usages. The item is recorded and left open.

9.7.5 The “Trojan Horse” Naming Collision

§9.1.6 records that two source documents use the term “Trojan Horse” for concepts of opposite valence, and that neither refers to the other. In p12r01 the term names an attack: tools presented as improvements to decision efficiency that carry a bias toward financial return, whose recommendations drift as their data degrades in a way the source describes as undetectable without a deliberately conducted audit. In the glossary pass, a “Trojan Horse” Cell is a sanctioned influence tactic: a deliberate placement of values-aligned members into ventures with the stated intent of knowledge-sharing and advocacy, aimed at influencing external partners, and explicitly disclaimed by that source as not subversive (`~/code/site-provide-coop/!! Localization.md` , term 9; the term does not appear in `Values--.txt` , which begins at term 11).

Chapter 09 uses the attack sense throughout, and the catalog entry at §9.1.2 and the treatment at §9.4.5 both carry that sense. The collision itself is unresolved.

§9.1.6 states the operational consequence and the reason the choice is not made there. A term naming both an attack vector and an approved practice cannot appear in an audit finding or a detection rule without a qualification attached each time it is used, so one of the two usages should be renamed. Which one should be renamed determines which body of existing material has to be revised, and no source in the extraction has standing over both. §9.1 also fixes the constraint that any renaming has to work within: threat names in this chapter are the sources' own, retained so that a later reader can trace an entry back to the material that produced it. Renaming the attack sense would break that traceability for one catalog entry and for the two sub-patterns filed under it at §9.4.6 and §9.4.7. Renaming the sanctioned-practice sense would leave the catalog intact and would instead require revision of the glossary material and of anything downstream of it. That comparison is this chapter's, and it is recorded as a description of the cost on each side rather than as a recommendation. The choice is left open.

9.7.6 Who Conducts the Mechanisms

Limit 4 of §9.6 states that the mechanisms cataloged there name no one to run them. The gap is recorded here in the specific form it takes in each case.

Several entries specify a body or a role and stop at that. §9.6.3 names a designated role that analyzes patterns in word usage across the network, and does not say who designates it. §9.6.2 is one line in a terminology list: it specifies standing scrutiny of the network's decision-support tools and states no procedure, no cadence, and no party responsible for conducting the audit. §9.6.1 extends audit coverage past financial anomalies to patterns of conduct without naming the auditor. §9.6.4 specifies analytical systems examining network data for patterns of factional alignment, and §9.6.17 specifies analysis that highlights unusual clusters of activity for human examination, with the source stating as part of the mechanism that the output is a starting point for an investigation rather than a substitute for one; neither entry says whose examination or whose investigation. §9.6.16 specifies a mandatory transparency rotation on the committees that allocate resources and public dashboards recording the justification given for a decision, without naming the party that administers either.

§9.6.6 is the partial exception. Its two forms specify how members of the parallel body are chosen, by lot, and the second form specifies strictly limited terms and veto standing. Selection by lot answers who sits on the body. The following is this chapter's reading: it does not answer who convenes the body, what supplies it with the information a comparison would require, or what obliges the official body to respond when the two diverge.

One source raises the question against its own proposal. The analytical systems in the first form of §9.6.4 are to be kept under constant scrutiny so that they do not become instruments of control in their own right. The source states the condition and names no mechanism for meeting it, so the proposal that most clearly identifies the problem also leaves it open.

The consequence for this chapter is bounded. Chapter 03 specifies affirmative governance design, including how bodies are composed and how authority over a decision is settled, and this chapter cross-references it rather than restating it. Nothing in the extraction connects any mechanism in §9.6 to that design, and no source assigns any of these mechanisms to a body the model already specifies. The mechanisms are recorded as the sources state them, with the appointment, funding, and accountability questions unanswered.

9.7.7 Whether the Scenario at §9.6.16 Warrants a Catalog Row

§9.6.16 records a mechanism the primary-source sweep recovered: a single prevention line combining a mandatory transparency rotation on resource committees, public dashboards recording decision justifications, and gamified challenges in which members analyze the network's own data for resource manipulation conducted in the form of oversight. The source attaches that line directly beneath a scenario of its own, which makes it one of five source-level attachments in the whole material. It is one of two that land on nothing §9.1.2 catalogs; the other is the onboarding prevention line at §9.6.19, whose scenario is a member presenting as a mentor to newly formed units and misrepresenting the network's policies to them. The scenario treated here is an internal coalition circulating accounts of bias in resource-allocation decisions, algorithmic and human, in order to portray the network's leadership as corrupt and to degrade trust in it.

The scenario is close to two cataloged entries without matching either. Manufacturing Dissent (§9.4.8) works by manipulating what the network's health-monitoring systems display, so that members respond to a displayed pattern the underlying records do not support; this scenario works by circulating accounts about decisions, and requires no manipulation of a system. Divide and Conquer (§9.4.4) works by circulating unverified accounts, which is the same mechanism, but the source assigns it to outside parties and directs it at trust between units; this scenario is conducted by an internal coalition and directed at the standing of the bodies that allocate resources. §9.6.16 records the link to Manufacturing Dissent as inferred for those reasons.

Recommendation. This chapter recommends that the scenario receive its own row in §9.1.2, and states this as a recommendation rather than as a change made here. The argument is the catalog's stated function: §9.1 exists so that a countermeasure can be traced back to what a source directed it at, and leaving this scenario uncataloged means one of the five source-stated threat-to-countermeasure links in the entire material has no anchor in the catalog, which is the same directedness deficit §9.7.8 records for the chapter as a whole. The row would belong in the external and technical vectors category alongside §9.4.4 and §9.4.8, and its name would have to come from the source's own heading text so that §9.1's naming rule holds; the source supplies a timeline heading and a target rather than a threat name, which is a weaker naming basis than the other 22 rows rest on.

The same argument stands for the onboarding scenario at §9.6.19, whose naming basis in its source is of the same kind and of the same weakness. It is not carried into the recommendation here, because the arithmetic below was worked for one added row against the catalog in its current form. A later editor weighing one of the two should weigh both.

The arithmetic consequences are set out so that a later editor can weigh them, and they are this chapter's. The catalog would hold 23 rows. The new row would be `addressed`, since its mechanism carries a source-level attachment, raising the addressed count from two to three. Manufacturing Dissent would then rest on §9.6.18 alone, which the sweep rates its least confident item and which §9.6.21 already states its status does not rest on, so that row would move from `partial` to `none proposed`. The resulting distribution is three addressed, thirteen partial, and seven with nothing proposed. The change is not made here because the counts recorded in §9.6.21 were verified against the catalog in its current form.

9.7.8 What the Coverage Table Records

The tables at §9.6.21 cover 22 threats and eight failure modes. Two threats are recorded as addressed, fourteen as partial, and six as having nothing proposed against them. No failure mode is recorded as addressed; seven are partial and one has nothing proposed.

Fourteen partial entries overstates how much of the material is directed at a named threat, and the cells in §9.6.21 state why one at a time. Exactly one of the fourteen rests on an attachment a source makes itself: the framework-capture scenario of §9.3, which the source presents as unresolved and answers with three responses it does not offer as a resolution. Six rest on correspondences this chapter or the primary-source sweep drew. The remaining seven rest on correspondences the staged extraction drew. A reader counting what the sources actually direct at a named threat should count the two addressed entries and that single partial, and should read the other thirteen as a record of what bears on a threat rather than of what anyone proposed for one.

The failure-mode table carries the same qualification in a different form. Its seven partial entries all rest on the instruments §9.6.20 assigns to chapter 03, which are general public-administration content rather than mechanisms specific to this model: the extraction says so of the material it carried, and the fourth polycentric response, recovered into that subsection, is of the same kind. No failure mode in §9.5 has a mechanism specific to this model directed at it.

The reason is a property of how the sources are written, and §9.6 limit 2 states it. In the source documents the countermeasures appear as lists attached to a response's overall subject rather than as a mapping onto named threats. Five attachments in the whole material are made at the source level, and two of those five land on scenarios the catalog does not name (§9.7.7). Every other correspondence in §9.6.21 was drawn afterward — in the staged extraction, in the primary-source sweep, or in this chapter — and each is marked in the cell where it appears.

The following statement of what that implies is this chapter's. A status in §9.6.21 answers the question of whether anything in the material bears on an entry. It does not answer whether anyone proposed a mechanism for that entry, and it does not answer whether the mechanism would work, since nothing in §9.6 has been built or tested. Those are three separate questions, and the table answers the first. Three of the 22 threats have an answer to the second — Mission Creep, Lexical Drift, and the framework-capture scenario, the three rows whose entries carry a source-level attachment. None has an answer to the third.

10.0 • Evaluation

The performance of a Humanized Autonomous Organization (HAO) — the network’s coordinating framework — depends on structural integrity and economic efficiency, and on its ongoing capacity to sense, reflect, and adapt in alignment with its foundational principles. This section describes how HAOs embed intelligence into their operations through distributed sensing, participatory evaluation, and recursive learning mechanisms that span individuals, teams, and the network.

Unlike traditional performance management systems, which often reduce organizational health to financial metrics and lagging indicators, HAOs use **multi-capital frameworks** to evaluate health across social, cultural, ecological, and intellectual domains. Evaluation is formative, relational, and reflexive rather than punitive: it supports trust, identifies misalignments, surfaces learning opportunities, and reinforces coherence across autonomy.

To this end, HAOs deploy a combination of human practices and technological systems — the **Collaborative Intelligence Network (CIN)** — to support visibility, accountability, and adaptation across the network.

This section introduces:

- The **multi-capital evaluation framework** used across HAOs
- Protocols for assessing participation quality, trust alignment, and role health
- Continuous sensing mechanisms embedded within UMEs (small, self-managing venture teams) and network operations
- The role of AI-augmented feedback systems in supporting distributed governance
- Systemic learning loops that drive the evolution of governance, agreements, and culture

Taken together, these elements make up the intelligence fabric of a HAO, supporting alignment with its values while navigating uncertainty, complexity, and growth.

§10.1 Metrics Across Capitals

Traditional organizations often rely on a narrow band of indicators, primarily financial metrics such as profit, growth rate, and return on investment, to evaluate performance. Humanized Autonomous Organizations (HAOs) — the network’s coordinating framework — use a multi-dimensional view of value creation and systemic health, tracking trust, autonomy, resilience, and shared meaning alongside financial returns. These qualities cannot be fully understood through financial indicators alone. HAOs must assess what is growing, how, for whom, and at what cost to long-term integrity.

To capture the complexity of distributed, value-aligned organizations, HAOs employ a **multi-capital evaluation framework** that spans five distinct yet interdependent domains: financial, social, cultural, ecological, and intellectual capital. This approach draws on systems thinking, treating each form of capital as co-produced and interdependent. Together, these dimensions offer a fuller picture of what the organization is producing, how it is evolving, and whether it remains aligned with its core principles.

Metrics across these domains serve as relational signals: tools for shared reflection, proactive correction, and adaptation, rather than mechanisms for centralized control. They can be surfaced through a combination of automated monitoring, peer review, self-reporting, and AI-augmented analysis, with an emphasis on transparency, agency, and ethical interpretation.

10.1.1 Financial Capital

Objective: Track the flow, efficiency, resilience, and fiscal sustainability of economic resources and capital flow coherence across the network.

Indicators:

- **Revenue Throughput:** Gross and net revenue at the UME (a small, self-managing venture team, $\leq \sim 15$ people), SEP (a joint venture between teams), and HAO levels, disaggregated by maturity stage.
- **Capital Allocation Efficiency:** % of initial investment reaching edge-level (UME) activity within thirty days of capital intake by the HAO.
- **Break-Even Velocity:** Median duration (in months) from UME launch to financial self-sufficiency — consistent revenue generation exceeding expenses.
- **Diminishing Contribution Ratio:** Ratio of upward revenue contributions from UMEs to HAO infrastructure against outward reinvestment, tracked over time to show diminishing contribution patterns.
- **Liquidity Resilience:** Operational runway (in months) per node — UME, SEP, and HAO reserve — assuming no incoming revenue.
- **Debt Load Index:** Internal borrowing vs. external liabilities per capital-receiving entity, as a measure of economic sovereignty.

Interpretation: These indicators show whether the economic mechanism of the HAO — its trickle-out flow architecture — supports its intended purpose: to direct investment to the edge, reward value creation, and sustain distributed operations without overburdening contributors, concentrating power, bottlenecking liquidity, or exposing the network to systemic fiscal stress.

10.1.2 Social Capital

Objective: Measure the strength, density, and quality of human relationships — trust, collaboration patterns, and perceived fairness — that enable cooperation and shared governance within and across HAOs.

Indicators:

- **Network Trust Index:** Composite trust score derived from regular sentiment-based survey assessments across the network, indexed to ETHICAL values.
- **Participation Density:** Proportion of members actively engaged in governance, working groups, feedback, mentorship, or SEP collaboration within the last 90 days.
- **Cross-UME Connectivity Rate:** Frequency and depth of collaboration between UMEs, including the number of active SEPs, shared functions, shared projects, and mutual aid.
- **Conflict Recovery Lag:** Average time between the surfacing of a breakdown in trust and its restoration via formal mechanisms.
- **Onboarding Integration Time:** Time from a new member’s entry to their first meaningful participation in governance or operational decisions.

Interpretation: Social capital is the connective tissue of a HAO. These indicators show how well the organization is fostering belonging, reciprocity, and distributed agency, and help identify emergent friction, disconnection, or exclusion patterns before they degrade network cohesion.

10.1.3 Cultural Capital

Objective: Evaluate the coherence and vitality of shared values, norms, language, and purpose within and across diverse or semi-autonomous UMEs.

Indicators:

- **Value Alignment Score:** Rate at which peer-reviewed decisions or behaviors align with the network’s ETHICAL principles, verified through VAM (ongoing checks that actions match stated principles).
- **Cultural Drift Rate:** Degree of divergence from previously agreed cultural norms or principles, tracked longitudinally.
- **Memetic Coherence:** Frequency and consistency of usage for shared concepts, metaphors, core terminology, and narrative references across discourse and communications.
- **Cultural Fluency Penetration:** Proportion of members proficient in using internal concepts such as “PEM,” “DEA” (a versioned operating agreement replacing fixed bylaws), and “SEP.”

Interpretation: These indicators function as early signals for fragmentation, ideological divergence, or misalignment. They quantify the shared meaning-making that enables coordination across distributed units, helping reinforce coherence without enforcing homogeneity.

10.1.4 Ecological or Planetary Capital *(Contextual/Optional)*

Objective: Where applicable, assess the environmental impact of UME and network operations — including resource replenishment, restoration, and alignment with the resource cycles of place-based systems.

Indicators:

- **Resource Intensity:** Water, energy, and material consumption per unit of output.
- **Community Impact Score:** Feedback from local ecosystems and community stakeholders on the presence and activity of UMEs (e.g., ecological, noise, and traffic impacts).
- **Circularity Index:** Percentage of materials and resources reused, recycled, regenerated, or cycled within the network.
- **Commons Integrity Rating:** Degree to which shared ecological or infrastructural resources (e.g., land, code, infrastructure) are maintained, enriched, or preserved under collective management.

Interpretation: For place-based or ecologically grounded implementations, these indicators help align operations with planetary boundaries and bioregional stewardship, tying stewardship indicators directly to business decisions.

10.1.5 Intellectual and Knowledge Capital

Objective: Track the system’s capacity to learn, adapt, innovate, and retain collective intelligence.

Indicators:

- **Protocol Evolution Velocity:** Frequency, depth, and number of substantive changes to governance, coordination, or economic structures per cycle.
- **Knowledge Redundancy Index:** Percentage of critical operational knowledge that is held by at least two individuals or systems.
- **Pattern Library Growth:** Accumulation of documented practices, principles, learnings, and actionable insights (“memes”) codified across UMEs and SEPs over time.

- **Institutional Memory Integrity:** Continuity of, and access to, historical knowledge and rationale for prior decisions (e.g., reason for DEA clauses), measured through metadata and traceability.
- **Open Contribution Rate:** Proportion of members who regularly contribute to the knowledge commons — documentation, code, governance improvements, patterns, or shared learning spaces.

Interpretation: These indicators support the antifragility of the organization. They reveal how knowledge is generated, shared, and preserved, especially across transitions and scaling, and indicate whether the organization can absorb the loss of any single node without systemic failure.

Summary

Each of the capital domains outlined in this section is essential to the overall health of the HAO. While financial metrics remain necessary for operational viability, they are insufficient on their own. Social and cultural metrics measure the integrity of human systems; ecological metrics track resource use and renewal; intellectual metrics track the learning capacity of the network itself.

By deploying these metrics in a distributed, non-coercive manner, augmented by participatory review and AI-supported pattern recognition, HAOs can maintain an ongoing view of their health and alignment and remain adaptable and aligned with their founding principles as they grow and evolve. These metrics support dialogue, reflection, and distributed decision-making in addition to diagnosis, functioning as tools for continuous alignment and systemic feedback and forming the basis of the intelligence and governance functions described in the sections that follow.

When integrated into participatory feedback loops, these metrics also function as relational assets that strengthen collective intelligence and shared responsibility.

§10.2 Participation Quality and Alignment Audits

Evaluation in a Humanized Autonomous Organization (HAO) — the network’s coordinating framework — considers quantitative engagement metrics, such as how many people voted or how often they showed up, together with the quality, depth, and systemic coherence of participation. Participation is not inherently valuable unless it is aligned with the system’s principles, contributes to shared learning, and strengthens the fabric of trust. In this sense, HAOs treat participation as both a right and a responsibility, grounded in autonomy but bounded by shared purpose.

To support this, HAOs implement structured but adaptive **participation audits** that assess how much people participate, how well, in what roles, and with what systemic effects. These audits are conducted through a combination of peer feedback, consent tracking, role health mapping, and governance analytics. These systems are designed to support participation through clarity, reciprocity, and feedback rather than coercion.

10.2.1 Consent Decision Audits

Consent-based governance (where decisions move forward unless there is a substantive objection) is a core method within many HAO frameworks. Tracking the integrity of consent processes helps ensure that this mechanism does not become performative or exclusionary.

Audit Criteria:

- **Consent Velocity:** Average time from proposal introduction to consent resolution.
- **Objection Validity Rate:** % of objections categorized as structural (e.g., violates principles), emotional (e.g., unresolved tension), or procedural (e.g., process violation).
- **Consent Coverage:** % of members or roles meaningfully informed about a proposal before it passes.
- **Objection Resolution Lag:** Time from raised objection to consensus or reframed proposal.
- **Silent Pass Rate:** % of proposals approved with no comments or objections; high rates may indicate disengagement.

Purpose: Ensures that “no objection” is not mistaken for alignment, and that objections are seen as valuable inputs rather than disruptions.

10.2.2 Role Health Assessments

In HAOs, individuals typically inhabit multiple roles (some elected, some emergent, some rotating). These roles form the functional scaffolding of the system, and their health is a direct indicator of systemic resilience.

Audit Criteria:

- **Role Fulfillment Variance:** Difference between defined responsibilities and actual activity performed within a role.
- **Role Overload Score:** % of members holding three or more high-stakes roles simultaneously.
- **Rotation Adherence:** Frequency of mandated or encouraged rotation cycles (especially in governance roles).
- **Role Density Index:** Ratio of active roles to active members, used to detect role inflation or under-participation.
- **Exit Continuity Health:** Existence of documented transitions and handoffs when roles change.

Purpose: Helps prevent burnout, overcentralization, and stagnation while fostering clarity and equitable distribution of responsibility.

10.2.3 Equity-Engagement Correlation Mapping

While many HAOs provide members with equity, governance power, or shared ownership, having a stake does not guarantee meaningful participation. Conversely, some highly engaged members may hold little formal power. This audit maps the relationship between value received and value contributed, identifying gaps, misalignments, or systemic vulnerabilities.

Audit Criteria:

- **Ownership-Activity Correlation:** Graphing participation frequency against equity shares.
- **Underleveraged Talent Index:** Identification of members with high potential or capacity but low current engagement.
- **Passive Extraction Risk:** Detection of actors who consistently receive network value (distributions, influence) without contribution.
- **Equity Diversification Score:** Concentration of governance or financial stake within top 5% of members.

Purpose: Supports balance between ownership and contribution, so that equity reflects ongoing system participation rather than a fixed entitlement.

10.2.4 Network-Level Engagement Distribution

To avoid centralization, HAOs must remain attentive to who is participating, from where, and in what ways. This involves visualizing and evaluating participation across roles, nodes, and time horizons.

Audit Criteria:

- **Engagement Diversity Index:** Degree of participation spread across different UMEs (small, self-managing venture teams, $\leq \sim 15$ people), geographies, or demographic groups.
- **Temporal Distribution Patterns:** Mapping participation spikes and lulls over time; useful for surfacing coordination fatigue.
- **Governance Participation Depth:** Average number of governance processes a member engages in per quarter.
- **SEP Engagement Ratios:** Number of SEPs (joint ventures between teams) a member or UME is connected to, as an indicator of embeddedness.
- **Feedback Reciprocity Rate:** Ratio of feedback given vs. received per member or team.

Purpose: Ensures that engagement is not clustered in a narrow subset of actors, that opportunities for involvement are widely available, and that systemic blind spots are surfaced early.

Summary

Participation in a HAO is evaluated by alignment, intention, and contribution to systemic coherence, not by presence alone. The audits described here assess engagement using more dimensions than a simple “active vs. inactive” classification.

By periodically surfacing these insights (via automated tools, governance reviews, or peer processes), HAOs can build environments where participation is visible, supported, and self-reinforcing. Rather than enforcing minimum standards, this architecture makes the health of the whole a shared responsibility that members take up voluntarily.

§10.3 Continuous Sensing and System Adaptation

In traditional organizations, sensing is often retrospective, tied to quarterly reports, static KPIs, or after-the-fact surveys. Humanized Autonomous Organizations (HAOs) — the network’s coordinating framework — require systems that are proactively perceptive, locally attuned, and continuously adaptive. Operating across complex, evolving environments, HAOs must detect emergent patterns, respond to early signals of misalignment, and evolve both their decisions and their decision-making processes.

Effective sensing requires designing for pattern recognition, relational awareness, and timely feedback at all levels of the system. Sensing within HAOs is distributed: each UME (a small, self-managing venture team, $\leq \sim 15$ people), SEP (a joint venture between teams), or role functions as a node in a distributed sensing network. This sensing structure is composed of three interlocking strata: internal feedback protocols, ambient monitoring systems, and localized cultural sensors.

Together, these systems support a shift from reactive management to ongoing monitoring and adjustment.

10.3.1 Internal Feedback Protocols

Each node in a HAO (whether a UME, working group, or governance circle) maintains structured feedback processes as part of its operational rhythm. These mechanisms create predictable opportunities for reflection, emotional processing, and constructive course correction.

Core Practices:

- **Cycle-Based Retrospectives:** Held monthly or quarterly, these sessions reflect on purpose, process, and people, as well as outcomes. Formats vary but often include check-ins, timelines, appreciations, tensions, and improvements.
- **Pulse Checks:** Brief, frequent surveys (weekly or bi-weekly) assessing clarity, morale, trust, and perceived coherence. Responses are anonymized and aggregated.
- **Consent-Based Feedback Loops:** Proposals or operational changes include built-in review periods, where members can offer feedback before and after implementation.
- **Listening Roles or Circles:** Certain members may hold explicit roles as listeners, confidants, or emotional anchors, ensuring that interpersonal signals are surfaced and processed.
- **Feedback-on-Feedback Routines:** Meta-level check-ins on how well the system is processing and responding to feedback itself.

Purpose: These internal feedback protocols are designed to turn governance and operational rhythms into learning loops, keeping human experience central to adaptation.

10.3.2 Ambient Monitoring Systems

Beyond explicit feedback, HAOs rely on ambient sensing mechanisms that track ongoing signals (financial, cultural, relational, and structural) without overburdening participants. This layer is primarily supported by automated systems and the Collaborative Intelligence Network (CIN).

Key Features:

- **Decision Latency Tracking:** Measures how long proposals, objections, and revisions take to resolve across the network.
- **Contractual Friction Index:** Monitors slowdowns, escalations, or repeated exceptions in dynamic agreements between nodes.
- **Participation Entropy Score:** Detects disproportionate drops in participation from individuals or teams over time.
- **Narrative Drift Detection:** Tracks changes in language, metaphors, or framing across public communications, useful for detecting slow ideological shifts.
- **Value Alignment Signals:** Aggregates decision metadata (e.g., objections raised, values cited) to flag divergence from declared principles, via VAM (ongoing checks that actions match stated principles).

Data Ethics Layer: Ambient sensing is opt-in and purpose-limited, and remains subordinate to human interpretation. Signals inform awareness, not dictate response. Raw behavioral data is never interpreted as intent.

Purpose: Ambient sensing complements active feedback with low-friction, low-noise insight, allowing the network to identify misalignments or decay patterns before they escalate into failure modes.

10.3.3 Localized Cultural Sensors

Culture is context-specific. While HAOs maintain shared values, they do not impose uniform norms. To stay coherent while remaining adaptive, each UME or SEP cultivates localized sensing mechanisms tuned to its own context.

Mechanisms:

- **Culture Sentinels:** Individuals or sub-groups tasked with noticing shifts in tone, trust, rituals, and alignment within their UME or SEP. Their role is limited to reflecting and reporting on these shifts.
- **Boundary Watch Processes:** Designed to detect early signs of cultural drift at the edge of a UME, especially in how it interfaces with other entities or new members.
- **Event Reflection Loops:** After moments of rupture (conflict, breakdown, or disagreement), the system invites structured reflection, both locally and network-wide.
- **Cultural Resonance Surveys:** Periodic qualitative instruments that ask members to describe their experience of purpose, alignment, and belonging using their own words.
- **Language Pattern Mapping:** CIN-supported analysis of recurring phrases, metaphors, or storytelling formats to detect resonance or dissonance with network-wide values.

Purpose: These cultural sensors help HAOs distinguish adaptation that preserves network coherence from fragmentation that breaks it, supporting both coherence and pluralism.

Summary

Sensing in HAOs is built into the design rather than added afterward. By embedding structured feedback, real-time pattern awareness, and localized cultural monitoring, the HAO can respond to both explicit inputs and subtle shifts across its human and technical systems. These mechanisms are intended to help the organization evolve with its members, course-correct without disruption, and maintain trust over time.

Continuous sensing is an ongoing organizational practice as well as a technical function, supporting the organization's ability to track and respond to change over time.

§10.4 AI-Augmented Governance Monitoring

As HAOs (the network's coordinating framework) operate across distributed teams, diverse cultural contexts, and evolving governance models, the scale and complexity of their internal systems can exceed what human cognition alone can manage. HAO governance design keeps final judgment with humans: artificial intelligence (AI) supports the people who govern rather than governing directly.

The role of AI in HAOs is to support collective awareness, aid pattern recognition, and assist with timely, principled decision-making. This is achieved through the Collaborative Intelligence Network (CIN): a suite of tools, protocols, and learning systems designed to help participants navigate complexity without depending on opaque technologies.

This section outlines how HAOs use AI to support governance monitoring across three layers: human-in-the-loop design, anomaly detection and foresight, and ethical alignment tracking.

10.4.1 Human-in-the-Loop Design

HAOs do not implement autonomous governance. AI systems provide **augmented agency**: they operate within clear constraints, providing insights, summaries, or flags. They do not execute decisions or enforce authority.

Design Features:

- **Proposal Summarization:** AI-generated distillations of governance proposals, surfaced alongside original text for faster comprehension and feedback.
- **Contradiction Detection:** Identification of proposals or decisions that conflict with existing value statements, recent commitments, or Dynamic Enterprise Agreements (DEA) — a versioned operating agreement replacing fixed bylaws.
- **Objection Clustering:** Semantic grouping of objections or concerns during governance cycles, helping facilitators identify patterns across responses.
- **Decision History Recall:** Timeline-based retrieval of relevant past decisions, enabling context-aware deliberation without manual search.
- **Smart Reminders:** AI nudges that prompt role holders or contributors when governance tasks are overdue or when required input is missing.

Safeguards:

- No AI system may execute governance changes, approve proposals, or override objections.
- All outputs are auditable and subject to participant review.
- "Pause and reflect" controls are available to disable or question AI-generated insights during deliberation.

Purpose: These tools reduce cognitive load and improve clarity. They do not replace the deliberative or ethical reasoning that is core to governance in an HAO.

10.4.2 Anomaly Detection and Scenario Simulation

AI systems in HAOs also detect early warning signs of misalignment and simulate the possible impacts of structural or economic decisions before implementation.

Anomaly Signals:

- **Centralization Drift:** Rising concentration of influence, equity, or decision-making within a small subset of roles or nodes.
- **Exit Clustering:** Higher-than-expected member or UME (a small, self-managing venture team, $\leq \sim 15$ people) departures, especially from specific circles, roles, or geographies.
- **Governance Fatigue Signals:** Declines in proposal quality, participation rate, or decision velocity over time.
- **SEP Imbalance:** SEP (a joint venture between teams) configurations that disproportionately favor one UME or exclude smaller players.

Simulation Scenarios:

- **Capital-Flow Disruption Modeling:** Predicts system-wide impact of delayed diminishing contributions, failed outward reinvestment cycles, or UME insolvencies.
- **Governance Change Simulations:** Projects how new protocols would affect participation, alignment, or trust metrics across stakeholder groups.
- **Crisis Modeling:** Runs stress tests based on historical shocks (e.g., UME collapse, conflict spikes, external market shifts).

Purpose: These mechanisms do not produce definitive answers. They provide structured foresight to help HAOs develop responses to emerging risks and opportunities.

10.4.3 Ethical Alignment Monitoring

One of the more sensitive uses of AI in HAOs is tracking the alignment between decisions and declared values. This function tracks coherence and drift; it does not impose punishment or conduct surveillance.

System Components:

- **Value Reference Index (VRI):** Embeds ETHICAL and PARTS principles into a semantic model against which governance activity is compared.
- **Alignment Flagging:** Detects when decisions, communications, or proposals significantly diverge from stated principles (e.g., lack of transparency, power hoarding).
- **Pattern Recognition:** Identifies recurring governance tensions, such as decisions made without meaningful consent or persistent bypassing of conflict resolution steps.
- **Impact Reflection Triggers:** Prompts human-led reviews when alignment violations occur consistently or involve high-stakes actors.

Consent Mechanisms:

- Members must consent to their governance activity being included in alignment modeling.
- AI outputs are always advisory; interpretation and response remain human responsibilities.
- Transparent logs allow any member to see how value judgments were derived or flagged.

Purpose: This layer helps the HAO track consistency between stated values and governance activity, without converting ethics into compliance checklists or bureaucratic rituals.

Summary

AI-augmented governance monitoring does not replace human decision-making; it supports pattern recognition and coordination across a distributed network. When designed with clear constraints, these systems help HAOs stay responsive to complexity while decisions remain grounded in human judgment, cultural plurality, and shared purpose.

§10.5 Systemic Learning and Evolution Protocols

The Humanized Autonomous Organization (HAO) — the network’s coordinating framework — is designed around deliberate, recursive learning. Unlike static institutions that treat policy and process as fixed, HAOs are designed to evolve continuously: socially, economically, and technologically. This evolution follows intentional **learning protocols** aligned with the organization’s values, scale, and distributed structure, rather than arising by chance or central control.

Where previous sections explored sensing and evaluation, this section outlines how HAOs translate what they learn into intentional change through governance adaptation, pattern dissemination, and capacity-building. Learning functions as infrastructure rather than an episodic activity. It operates at individual, team, UME (a small, self-managing venture team, $\leq \sim 15$ people), and network-wide levels, so that governance responds to feedback and continues to improve over time.

10.5.1 Learning Contracts for UMEs

Every United Micro Enterprise (UME) within the HAO is encouraged (or required, depending on maturity and funding stage) to operate with an evolving **learning contract**: a shared commitment to growth, reflection, and contribution to collective intelligence.

Core Components:

- **Quarterly Learning Objectives:** Co-developed by the UME with support from the HAO or peers; focused on improving internal processes, aligning values, or solving system-relevant challenges.
- **Reflection Cadence:** A regular rhythm (e.g., monthly reflection loops or retrospectives) with outputs shared across the network.
- **Peer Learning Exchanges:** At least once per cycle, each UME hosts or attends a session to share insights, tools, or failures with another unit.
- **Learning-Linked Benefits:** Access to capital, strategic support, or governance privileges may be linked to the consistent fulfillment of learning commitments.

Purpose: Learning is treated as a core output of every economic and organizational activity, not as a side effect.

10.5.2 Governance Reflection Cycles

HAOs maintain versioned constitutional frameworks (e.g., the Dynamic Enterprise Agreement) designed to evolve. Governance reflection cycles provide a structured opportunity to assess and adapt those frameworks based on experience and emerging needs.

Key Practices:

- **Annual or Biannual Reflection Forums:** Network-wide or federated gatherings to review the performance of governance protocols.

- **Multi-Capital Health Review:** Data from financial, social, cultural, and knowledge capitals are reviewed to inform updates.
- **Tension Mapping:** Aggregated lists of persistent systemic tensions, unresolved objections, or friction patterns.
- **Proposed Protocol Updates:** Participants offer structured proposals to evolve governance rules, decision paths, or participation mechanisms.
- **Consent-Based Ratification:** New versions of agreements are adopted via network-wide consent or representative mechanisms.

Purpose: Governance frameworks are designed to evolve without destabilizing the organization's core identity.

10.5.3 Pattern Library Development and Dissemination

The HAO maintains a library of organizational patterns: documented learnings, practices, rituals, and archetypes that emerge from across the network. These are context-aware, modular insights that can be adapted and reused, not fixed best practices.

Pattern Types:

- **Structural Patterns:** E.g., rotating facilitator models, shared resource hubs, fractal governance layers.
- **Process Patterns:** E.g., peer onboarding, rapid consent loops, feedback rituals, failure retrospectives.
- **Cultural Patterns:** E.g., storytelling formats, boundary-setting agreements, micro-trust repair mechanisms.

Library Architecture:

- **Decentralized Contributions:** Any member or UME may submit patterns using a standard template.
- **Versioning and Attribution:** Patterns evolve over time, with lineage and authoring preserved.
- **Cross-Context Annotations:** Each pattern includes notes on where it works, when it breaks, and how it adapts to scale.

Diffusion Mechanisms:

- **Inter-UME Learning Circles:** Curated conversations or jams where patterns are shared and refined.
- **Integrated Design Tools:** Governance platforms offer "pattern suggestions" during proposal drafting.
- **Public Commons:** Select patterns (redacted for privacy) may be shared publicly under open licenses.

Purpose: Builds institutional memory and coherence while preserving diversity across the network.

10.5.4 Adaptive Capability Investment

Learning requires resourcing to have an effect. HAOs invest time, attention, and capital into adaptive capabilities: addressing past problems and preparing for emerging needs and scenarios.

Core Approaches:

- **Learning Funds:** Pools allocated for capacity-building, training, and strategic adaptation.
- **Scenario Labs:** Short-term working groups tasked with modeling future states and prototyping governance or economic shifts.
- **Cross-Pollination Retreats:** Gatherings where members from different roles, UMEs, and geographies exchange insights and co-create new pathways.
- **Skill Stewardship Roles:** Designated participants who track emergent needs and mobilize internal talent to meet them.

Purpose: Embeds strategic foresight and capability development into ongoing organizational operations.

Summary

Legacy institutions often adapt reactively and after the fact. HAOs are designed to create conditions for ongoing, proactive, collective, and value-aligned learning. The learning protocols outlined here reinforce system resilience and are a mechanism through which HAOs maintain coherence as complexity increases.

Systemic learning functions as part of governance and organizational identity, not as a separate or optional activity.

11.0 • Deployment

§11.1 Seeding an HAO Network

Initiating a Humanized Autonomous Organization (HAO) — the network’s coordinating framework — depends less on capital or technology than on clarity of intent, defined architectural choices, and trust among founding participants. This section defines the minimum viable conditions, design parameters, and early actions required to instantiate a coherent HAO network.

11.1.1 Prerequisites for Genesis

Before operationalization, the initiating team defines several preconditions for the HAO model:

- **Shared Ethical Grounding:** All founding members commit to applying the ETHICAL framework (Empathy, Transparency, Harmony, Integrity, Collaboration, Accountability, Learning & Longevity) as constraints on design and decision-making.
- **Intent Declaration:** A *Statement of Purpose* is codified, answering:
 - What domain or problem space the HAO exists to support
 - Who the first beneficiaries are (members, communities, ecosystems)
 - Why distributed autonomy is more appropriate than centralized control
- **Initial Trust Fabric:** A minimum viable trust layer (formed through prior collaboration, trusted referrals, or progressive trust verification) is required for high-stakes cooperation under uncertainty.
- **Resource Alignment:** Founders must contribute one or more of:
 - Time-based capital (operational labor)
 - Knowledge capital (design, legal, technical)
 - Financial capital (seed investment, grants, mutual credit commitments)

These are tracked as early “risk contributions” and will later be converted to equity or influence under the HAO’s dynamic value model.

11.1.2 Selection of Domain Scope

While HAO is a general-purpose architectural pattern, early instantiation benefits from domain specificity to constrain complexity. Founders must choose:

- **Generalist HAO** (e.g., provide.io, focused on enabling infrastructure and services across sectors)
Pros: Broad applicability, supports many UMEs (small, self-managing venture teams of up to ~15 people)
Cons: Higher coordination load, delayed domain feedback
- **Domain-Specific HAO** (e.g., ICN (the reference cooperative business network) for cooperative enterprise, MTU (the network’s credit-union-like financial institution) for community banking)
Pros: Immediate feedback from context-specific needs
Cons: Requires deeper subject matter expertise

Selection of scope should consider:

- Level of regulatory complexity (e.g., healthcare vs. creative cooperatives)
 - Capital intensity and infrastructure readiness
 - Potential for UMEs to emerge organically within the chosen domain
-

11.1.3 Initial Governance Stack

The governance structure must be instantiated at the outset with minimal viable complexity but long-term extensibility. It should include:

Constitutional Layer

- Codified in a **Dynamic Enterprise Agreement (DEA)** v0.1 — a versioned operating agreement replacing fixed bylaws
- Specifies:
 - Core principles and right-to-exit
 - Value accounting primitives (contribution → equity → influence)
 - Process for proposing, accepting, or amending changes (versioned, not overwritten)

Strategic Layer

- Defines:
 - Short-term funding priorities
 - Network formation sequencing (e.g., which UMEs are bootstrapped first)
 - Role designations (e.g., conveners, facilitators, architects, fiduciaries)

Operational Layer

- Defines:
 - Decision-making mechanics (e.g., sociocratic consent, dynamic roles)
 - Conflict resolution protocol
 - Feedback cadence and audit schedules

This stack should be bootstrapped using open, forkable governance tools (e.g., DAO-style frameworks, sociocratic role systems, collaborative docs with versioning), with an emphasis on human readability and multi-language accessibility.

11.1.4 Role Formation and Scaffolding

Initial participants act as scaffolders for the broader system rather than as formal “members” of a UME. Suggested minimum roles include:

Role	Purpose
Convenor(s)	Steward initial alignment, facilitate purpose refinement
Architect(s)	Translate HAO concepts into operational, legal, and technical patterns
Legal Navigator	Design entity formation, agreements, and jurisdictional protections
Systems Integrator	Build the initial digital stack (ledger, schema, governance tools)
Cultural Anchor	Safeguard psychological safety, coherence, and value alignment

These roles are temporary and rotating, governed by the principle of *subsidiarity*: authority is held only as long as it is needed, and as close to the point of action as possible.

11.1.5 Readiness Gates and Activation

Before moving from prefiguration to operational launch, the founding team conducts a readiness assessment:

- Has the purpose been codified in shared language?
- Do all key roles have redundancy or succession plans?
- Has the DEA been ratified by all initial participants?
- Have initial capital contributions (time, knowledge, financial) been logged and committed?
- Has a minimum viable infrastructure stack been deployed and tested?

Once affirmative, the HAO is considered in “genesis mode” and proceeds to instantiate its first UME(s) and SEP(s) — joint ventures between teams — evolving through the full stack described in Section 11.2 onward.

§11.2 provide.io as a Deployment Engine for HAO

While the HAO (the network’s coordinating framework) is an architecture rather than a product, its operationalization requires an implementation entity: an organization that translates governance and value-flow principles into deployable systems. *provide.io* serves this role. It is a for-profit cooperative that builds, maintains, and evolves the infrastructural backbone of HAOs, aligned with the network’s human-centered design principles.

This section outlines how *provide.io* functions as a modular, domain-agnostic systems integrator, and how its operational patterns can be replicated, forked, or federated by other entities.

11.2.1 Systems Integration Role

The primary function of *provide.io* is to synthesize technological, organizational, legal, and cultural layers into a unified HAO operating environment. It acts as a cross-domain scaffolding team, with responsibilities including:

- Technical Deployment
 - Implementing foundational systems: Distributed Ledger Infrastructure (DLI), Collaborative Intelligence Network (CIN), Enterprise Culture Cultivation (ECC), and Value Alignment Monitoring (VAM) — ongoing checks that actions match stated principles

- Maintaining schema registries, identity layers, and consensus engines
 - **Governance Tooling**
 - Authoring and maintaining open-source templates for Dynamic Enterprise Agreements (DEAs) — versioned operating agreements replacing fixed bylaws
 - Building governance interfaces for multi-role access, versioned policy enforcement, and participatory decision-making
 - **Economic Coordination**
 - Designing and deploying contribution tracking systems (slice models, equity converters)
 - Managing temporary holding structures (e.g., shared caches, seed pools) for early network liquidity
 - **Legal Design**
 - Constructing modular legal wrapper templates (e.g., cooperative LLC hybrids, SEP (a joint venture between teams) legal scaffolds)
 - Providing jurisdictional analysis and fallback protocols
-

11.2.2 Organizational Architecture of provide.io

provide.io is structured as a for-profit, member-cooperative hybrid designed to reflect the values and principles of the HAOs it supports. Its internal configuration includes:

Layer	Description
Governance	Consent-based council system split between architecture, integration, and ethics
Legal Entity	Flexible Benefit LLC or Platform Cooperative with dynamic ownership ledger
Ownership	Dynamic equity split between founders, contributors, and a reserve held for downstream HAO participants
Revenue Model	Blended: infrastructure-as-a-service fees, equity positions in supported HAOs, and optional licensing of components
Exit Protection	Contractual limits on outside equity stakes, restricting acquisition of majority control by external investors

It remains non-central, maintaining majority control only where necessary to preserve integrity during early network formation. Post-stabilization, it supports decentralization via handoff, licensing, and protocol federation.

11.2.3 Infrastructure Bootstrapping Toolkit

The provide.io stack includes both technical modules and human-operational protocols, intended to instantiate a minimal viable HAO ecosystem. Core components include:

A. Technical Infrastructure

- **Distributed Ledger Infrastructure (DLI)**: Tracks transactions, contributions, and agreements
- **Schema Registry**: Formal definition of governance roles, capital flows, trust scores
- **CIN Layer**: Human-AI decision augmentation tools for proposals, dispute analysis, and adaptive modeling
- **Credential and Trust Framework**: Multi-layered identity and verification system (used by MTU (the network’s credit-union-like financial institution) and ICN (the reference cooperative business network) alike)

B. Operational Templates

- **DEA v0.1 Reference Blueprint**
- **UME Agreement Generator**: Generates the agreement for a UME (a small, self-managing venture team, $\leq \sim 15$ people), including revenue split, exit protocols, and cultural onboarding flows
- **SEP Co-Governance Templates**
- **Role Ledger + Contribution Tracker**: For time, knowledge, and capital inputs

C. Observability and Simulation Tools

- Simulated governance games (e.g., “soft votes” with replayable outcomes)
- Stress-testing protocols for economic resilience (using historical failure pattern libraries)

All components are modular, interoperable, and forkable, supporting other HAO-implementing entities that evolve independently.

11.2.4 Deployment Patterns and Forkability

provide.io does not hold exclusive rights to HAO deployment. Its structure and output allow other, similar entities to replicate its function in different regions or domains.

A. Patterns Supported

- **Direct Deployment:** provide.io sets up core systems for a new HAO (e.g., ICN Alpha)
- **Platform-as-Protocol:** provide.io only deploys the base layer (e.g., DLI + governance stack), leaving economics to domain actors
- **Handoff Pattern:** After bootstrapping, provide.io dissolves its role or transitions to a minority contributor
- **Federated Deployment:** provide.io collaborates with other system integrators to co-deploy infrastructure (useful in global or bioregional rollouts)

B. Federation, Not Centralization

- Cross-provider protocol interoperability via a **Common Protocol Interface (CPI Layer)**
- Shared registries and licensing frameworks
- Guardrails to prevent network capture: open-source requirements, non-transferable voting rights, and rotating control nodes

11.2.5 Case Reference: provide.io + ICN Instantiation

In the case of the Integrated Cooperative Network (ICN), provide.io serves as:

- Architect and maintainer of the initial Dynamic Enterprise Agreement
- Builder of early shared infrastructure (e.g., SEP stack, shared cache wallet, onboarding flow)
- Host of internal tools for slice tracking, revenue reconciliation, and cultural protocol alignment
- Deployment partner for the ICN’s first three UMEs, ensuring compliance with versioned governance

This serves as a reference model for future HAO initiators: a systems-integration-first approach, paired with a human-centered ethos, can instantiate complex decentralized networks without sacrificing coherence or resilience.

§11.3 Genesis Process for the First UMEs

United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people — are the foundational productive units of a Humanized Autonomous Organization (HAO) — the network’s coordinating framework. Their formation is both a social event and a systems-activation process. This section outlines how to instantiate the first cohort of UMEs within a newly formed HAO, aligned with the network’s values, economic logic, and governance structure from the outset.

11.3.1 UME Formation Criteria

Not every idea or group qualifies as a UME within a HAO. Early UMEs must be selected and cultivated based on their ability to serve as:

- **Proof-of-Pattern Entities:** Validating core elements of the HAO such as participatory governance, trickle-out economics, and dynamic agreements
- **Strategic Anchors:** Filling necessary capabilities in the ecosystem (e.g., technology, learning, logistics, finance)
- **Cultural Beacons:** Demonstrating how values are enacted operationally

Selection Criteria:

Dimension	Example Evaluation Criteria
Alignment	Clarity of purpose matching HAO mission and ethics
Cohesion	Pre-existing trust or shared context among founding members
Feasibility	Demonstrated ability to deliver a value proposition in <6 months
Differentiation	Fills a distinct role in the emerging SEP graph
Contribution Model	Clear contribution structure (time, capital, knowledge) mapped to slices

11.3.2 UME Setup Toolkit

To streamline UME instantiation, the HAO (via an entity like provide.io) offers a standardized **UME Setup Toolkit**, consisting of the following components:

A. Structural Templates

- **Standard UME Agreement Template (SUAT)**
 - Defines purpose, role boundaries, revenue flow, exit terms
 - Modular: allows adaptation for service vs product UMEs, or solo vs collective structures
- **Contribution Ledger + Role Matrix**

- Tracks time, capital, IP, and relational contributions
- Roles mapped to compensation and influence pathways

B. Governance Onboarding

- Consent-based decision model with optional sociocratic scaffolding
- Onboarding to the HAO-wide Dynamic Enterprise Agreement (DEA) — a versioned operating agreement replacing fixed bylaws — including understanding:
 - Network obligations
 - Performance expectations
 - Participation in value-alignment monitoring

C. Infrastructure Provisioning

- Initial access to:
 - Shared ledger systems
 - Governance dashboard
 - Economic modeling tools (slice calculators, equity sim tools)
 - Cross-UME collaboration platforms (SEP (a joint venture between teams) management tools)

11.3.3 Soft Launch Protocol

Before full integration, each UME undergoes a soft launch cycle within a controlled operating envelope:

Stage 1: Intent to Form

- Founders sign a **Letter of Intent** to create a UME under HAO rules
- Initial role agreements and contribution estimates captured
- Submitted to the HAO strategic node for provisional approval

Stage 2: Provisional Operations (3–6 weeks)

- UME begins operation with limited autonomy
- All governance decisions logged and tagged for reflection
- Performance is tracked across three axes:
 - Delivery efficacy (Did it produce?)
 - Cultural fit (How well did it align?)
 - Relational feedback (Did it maintain psychological safety?)

Stage 3: Review + Integration

- A review process (can include peer UME members, HAO anchors, cultural stewards) determines:
 - Whether the UME graduates to full status
 - Whether structural changes are needed (e.g., role realignment)
 - Whether to merge, dissolve, or convert into a SEP instead

Upon successful graduation, the UME receives:

- Permanent ledger entries for its contribution history
- Access to full HAO benefits (profit share, governance rights, reinvestment channels)
- Rights to incubate or mentor additional UMEs in the future

11.3.4 Risk Containment and Boundary Management

Early-stage UMEs present financial, reputational, and structural risks. Mitigations include:

- **Operational Sandboxing**
 - UMEs operate in a contained scope with network access restrictions
 - Economic activities are tagged to “test-net” vs “main-net” value flows

- **Conflict Protocol Activation**
 - If significant friction arises, escalation to HAO mediators or cultural stewards
 - Optional use of “pause protocol” (temporary freeze on governance powers)
- **Exit Pathways**
 - Pre-negotiated paths for graceful dissolution, conversion to worker-owned business, or HAO-incubated recovery effort

11.3.5 Culture as Compliance

Each UME maintains operational, economic, and cultural coherence. HAO cultural protocols require:

- **Weekly Pulse Check**
 - Structured self-assessment of interpersonal trust, alignment, and emotional bandwidth
 - Logged to the Value Alignment Monitoring (VAM) system — ongoing checks that actions match stated principles
- **Narrative Logging**
 - Regular story-based reporting: “What did we try? What did we learn?”
 - Encourages sense-making and knowledge continuity
- **Cultural Failures as Signals**
 - Cultural degradation (gossip, stagnation, burnout) is treated as data, not pathology
 - Triggers support intervention, not punishment

This phased, values-anchored UME formation model is intended to let local autonomy operate without compromising network resilience or coherence.

§11.4 Strategic Enterprise Partnership Structuring

Strategic Enterprise Partnerships (SEPs) — joint ventures between teams — are the connective layer between United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people — within a Humanized Autonomous Organization (HAO) — the network’s coordinating framework. SEPs are formed to pursue shared objectives that exceed the capacity of a single UME, coordinating collaboration across UMEs in contexts requiring shared infrastructure, joint governance, or pooled resources. Unlike bilateral contracts or informal collaborations, SEPs are designed to follow the same governance principles as their constituent UMEs, while supporting domain-specific innovation and joint value creation. They support network-wide coordination without centralization: they connect UMEs through jointly governed, co-financed, value-generating ventures. This section presents a structured approach for forming, governing, and evolving SEPs within an HAO.

11.4.1 Purpose and Role of SEPs

SEPs exist to:

- Enable collective agency across multiple UMEs for goals requiring shared effort (e.g., logistics networks, digital platforms, co-branded ventures)
- Distribute risk, ownership, and decision-making across partners without requiring hierarchical control
- Formalize interoperability between UME governance structures
- Act as internal scaling mechanisms that preserve autonomy at the edge while achieving network-wide goals

A SEP is a governance container, economic contract, and operational unit that forms when more than one UME needs to collaborate at scale and with accountability.

11.4.2 Criteria for SEP Formation

SEPs are formalized, governed entities, instantiated when the following preconditions are met:

Dimension	Minimum Viable Condition
Purpose Scope	A goal or market opportunity that spans two or more UMEs
Strategic Alignment	The effort advances the purpose of multiple UMEs or the HAO at large
Contribution	Each UME brings defined value (labor, IP, funds, infrastructure); collaboration requires shared infrastructure, coordinated labor, pooled funding, or collective IP
Shared Risk	Outcomes, costs, and benefits affect all parties materially; potential downsides are distributed, and joint accountability is needed
Governance Need	Requires joint decision-making with formal accountability mechanisms
Non-Redundancy	The function cannot be adequately executed through isolated or informal agreements, or through one-sided subcontracting

Typical SEP domains and examples include:

- **Shared logistics networks**
 - **Distributed R&D labs:** collectives pooling IP development and talent
 - **Digital platform development:** joint service ventures in which multiple UMEs launch a shared platform or consultancy
 - **Market expansion strategies:** market-facing interfaces, e.g., a regional sales/distribution entity
 - **Regional or sectoral cooperatives**
 - **Public market buffer layers:** e.g., Contribulo-style entities mediating investor flows
-

11.4.3 SEP Governance Design

Each SEP operates under its own sub-governance protocol, derived from the Dynamic Enterprise Agreement (DEA) — a versioned operating agreement replacing fixed bylaws — but customized to the partnership’s needs. Each SEP is instantiated with a **SEP Charter** (also called a **Co-Governance Charter**), co-authored and co-signed by the founding UMEs and reviewed by the HAO governance anchor, the HAO Coordination Node, or an equivalent integrator (e.g., provide.io).

Governance Components:

1. **Charter:** Includes purpose, scope, expected duration, and founding contributions
 - Defines roles and decision mechanisms (e.g. rotating stewardship, quorum thresholds)
 - Establishes escalation pathways, exit procedures, and integration with HAO-wide Value Alignment Monitoring (VAM) protocols — ongoing checks that actions match stated principles
 2. **Governance Model:**
 - **Role-Based or Rotating:** Stewardship rotates or roles are functionally assigned
 - **Decision Protocol:** Consent-based for operational decisions; consent-based or supermajority for major changes; double-consent (all members + HAO coordination rep) for capital allocation or role changes
 - **Emergency Override:** Clauses invocable for financial or reputational risk events
 - **SEP Council:** Cross-UME group of role-holders acting as a governing forum
 3. **Conflict Protocols:**
 - Minor disagreements mediated internally
 - Structural disputes escalated to HAO-wide mediation or arbitration process
 4. **Exit Provisions:**
 - What happens if one UME departs
 - Reallocation of equity, IP, and commitments
 - Emergency clause if SEP endangers cultural or financial integrity of the HAO
-

11.4.4 SEP Economic Structuring

A SEP’s economic design addresses contribution accounting, revenue distribution, asset ownership, and equity/exit terms:

A. Contribution Accounting:

- All forms of capital tracked at UME level (time, financial, knowledge/IP, infrastructure, brand), covering both initial and ongoing inputs
- Use of contribution-ledgers or slice-models to capture dynamic value flows
- Slices assigned to participating UMEs, not individuals

B. Revenue Distribution:

- **Default Model:** Distribution proportional to initial and ongoing contribution ratios
- **Rebalancing Protocol:** Quarterly or milestone-based reassessment of ratios; future revenue splits revised via quorum-based reweighting, not hard-coded
- **Network Contribution:** Optionally, a network royalty (typically 5–15% — for example, 10%) flows to the HAO coordination pool

C. Asset Ownership:

- IP ownership jointly held by the SEP or conditionally licensed to members
- Access rights and sublicensing tied to participation or maturity of the partnership
- Commons licenses recommended for shared documentation and process artifacts

D. Equity and Exit Terms:

- SEP equity accrues to UMEs as organizations, not individuals
- Individual contributions rewarded inside their respective UMEs
- Buyout, sunset, or pivot scenarios documented in the Charter
- Defined dissolution protocol: IP handling, asset liquidation, equity rebalance across parent UMEs

11.4.5 SEP Lifecycle

To maintain clarity and continuity, SEPs follow a four-stage lifecycle model, each stage with associated protocols:

1. Initiation

- Joint proposal submitted by ≥2 UMEs
- SEP Charter + Economic Agreement drafted
- Provisional ratification by a neutral third party (e.g. HAO’s SEP committee); the HAO Coordination Node reviews for alignment and risk

2. Activation

- SEP goes live and enters operational mode, receiving provisional access to shared HAO tools, infrastructure, ledger, and reputation systems
- Onboarding of shared staff; monthly operating cadence begins: reporting, consent checks, and task distribution, with defined communication cadences
- Establishment of conflict protocol, activated at first dispute (to test robustness), and of a contribution audit cycle (typically quarterly)

3. Evaluation

- After a fixed cycle (e.g. 3–6 months post-launch), a formal, structured review of:
 - Outcomes vs. intent
 - Value creation vs. opportunity cost
 - Cultural coherence and feedback from participating UMEs
 - Operational effectiveness
 - Economic fairness
- Trigger conditions for revision:
 - Major UME exiting, or disengagement by a UME
 - Revenue threshold crossed, or revenue/cost imbalance
 - Legal or regulatory shift
 - Cross-boundary conflicts (e.g., resource hoarding, burnout)

4. Transition

- Post-evaluation paths:
 - **Stabilization:** SEP gains full member status and long-term continuity plan
 - **Handoff / Conversion:** SEP spins out as a standalone UME, affiliate organization, or multi-member platform coop
 - **Expansion:** SEP invites new UMEs to participate or replicates regionally
 - **Replication:** SEP becomes a reference pattern for other domains
 - **Dissolution:** SEP dissolves, and assets, obligations, value, and learning are redistributed

11.4.6 Integration with HAO-Wide Systems

Although SEPs are semi-autonomous, they remain legible and accountable to the HAO ecosystem, retaining interfaces with the broader HAO framework for integrity and adaptability.

Domain	Integration Mechanism
Governance Auditing	SEP governance and decision logs submitted quarterly to the DEA Archive or VAM registry
Trust Feedback	SEP behaviors influence cross-network reputation and alignment, affecting the trust scores of participating UMEs
Financial Compliance	SEP-related funding adheres to network-wide reinvestment rules and revenue transparency
Conflict Resolution	Major disputes may invoke HAO-wide mediation protocols
Observability	Performance data accessible via shared dashboards

Federation-Ready: A SEP pattern that works in one HAO can be exported to another — for example, from the ICN (the reference cooperative business network) to the MTU (the network’s credit-union-like financial institution) — via shared protocols and schema portability.

In some cases, a SEP becomes a precursor to a domain-specific HAO if its structure is self-sustaining and generalizable. For example, a logistics-focused SEP could evolve into a supply-chain HAO.

11.4.7 SEP Tooling and Templates

provide.io and other integrators offer:

- **SEP Formation Kits:** Charters, economic modeling tools, onboarding sequences
- **SEP Charter Builder:** Guided form for purpose, governance, and economics
- **Joint Contribution Ledger:** Per-SEP instance of slice/accounting infrastructure, tracking SEP-specific inputs across multiple UMEs
- **Consent Tracker:** Logs governance proposals, objections, and voting patterns
- **Conflict Simulation Tool:** Run-through of common SEP failure modes and resolution routes
- **SEP Ontology Registry:** Shared language for describing cross-UME roles and assets
- **SEP Replay Journal:** Stores contextual narratives and performance histories for future SEPs to learn from

Hosted and maintained by entities like provide.io or successor integration platforms.

11.4.8 SEP as a Bridge to Public Markets

Some SEPs may function as public interfaces, particularly in domains involving capital markets, external clients, or third-party licensing.

- **Example:** Contribulo-style buffer entity translating between traditional investors and HAO’s trickle-out economics
- **Firewall Design:** SEP maintains economic integrity while honoring external expectations
- **Governance Clause:** External board observers allowed; no voting control
- **Payout Models:** Revenue-share or capped returns, with clawback or community reinvestment triggers

This pattern lets external-facing entities participate in the ICN or HAO while limiting external control over internal decisions.

SEPs are formal, interoperable, accountable structures within the HAO. They let HAOs coordinate complexity without centralizing power, adding a federated economic layer for cross-UME collaboration and allowing coordination across UMEs without a central command structure. As the network matures, SEPs function as the mid-range connective layer between UMEs and HAO-wide coordination.

§11.5 Iterative Scaling Phases

Humanized Autonomous Organizations (HAOs) — the network’s coordinating framework — scale through iteration, coherence, and recomposition rather than top-down growth or central accumulation. This section outlines a phased model for scaling HAO deployments, based on feedback, modularity, and staged handoff rather than linear growth.

11.5.1 Phase One – Genesis Cluster (HAO + 3–5 UMEs)

This phase focuses on establishing trust, rhythm, and viability. It creates a dense social and operational nucleus around the HAO coordination function.

Key Objectives:

- Launch the HAO governance layer with minimal viable scaffolding
- Instantiate three to five diverse UMEs (small, self-managing venture teams) across complementary domains
- Deploy the initial economic and trust infrastructure (e.g. slice tracking, shared ledger)
- Establish first SEPs (joint ventures between teams) to test collaborative functionality

Indicators of Readiness for Next Phase:

- All UMEs have entered post-evaluation maturity
- First SEP has completed a full cycle (initiation, execution, review)
- Cultural trust protocol in place and functioning (e.g. pulse checks, reflection loops)
- Early members report “coherence and divergence”: shared purpose, diverse roles

11.5.2 Phase Two – Federation Layer Emergence (5–15 UMEs)

The HAO network now supports diverse, semi-autonomous value centers. This phase decentralizes governance and increases surface area for experimentation.

Key Objectives:

- Form at least two new SEPs between non-founding UMEs
- Distribute key HAO functions (e.g. onboarding, dispute resolution) to rotating or elected roles
- Launch UME-led governance experiments within boundaries of the Dynamic Enterprise Agreement
- Expand the cultural protocol to account for growing linguistic, operational, and emotional diversity

Infrastructure Additions:

- SEP registry
- Culture reflection archive or story-map system
- Role-based access control or verifiable credentials for governance authority

Risks:

- Divergence without cohesion (too much local experimentation without shared narrative)
- Coordination drag
- Burnout among original contributors

Scaling Milestone:

The HAO governance function becomes one among several centers of decision-making, rather than the sole organizing entity.

11.5.3 Phase Three – Network Topology Transformation (15–50 UMEs)

This phase transforms the HAO into a topological federation. Multiple governance nodes, economic flows, and innovation clusters emerge. The network behaves less like a star topology and more like a mesh.

Key Objectives:

- Enable interoperability across different types of UMEs (e.g. service, research, market-facing)
- Allow UMEs or SEPs to instantiate sub-HAO governance layers (nested or domain-specific)
- Launch cross-network initiatives (e.g. open toolchains, mutual credit systems)
- Form alliances or interoperability agreements with other HAOs or aligned networks

Infrastructure Additions:

- Inter-UME API schema registry
- Embedded simulation systems for governance or funding scenario modeling
- HAO-to-HAO trust bridge protocols (e.g. credential validation, shared investment mechanisms)

Outcomes:

- Functional redundancy across core services (no single point of failure)
 - Multiple cultural protocols coexisting under a shared values backbone
 - Cross-network arbitrage of knowledge, capital, and roles
-

11.5.4 Phase Four – Interoperable Network-of-Networks (50+ UMEs)

The system now operates as a constellation of federated governance systems, often spanning sectors, languages, or bioregions. Each HAO instance remains sovereign; interoperability replaces standardization as the scaling mechanism.

Key Objectives:

- Facilitate emergence of domain-specific HAOs (e.g. a supply-chain HAO, a learning HAO)

- Ensure interoperability through shared schema, legal templates, and data ontologies
- Formalize inter-HAO treaties or compacts to govern cross-network economic or legal activity
- Support migration or dual-membership across networks

Risks:

- Capture by capital (external or internal actors seeking control)
- Fragmentation into silos (loss of sense-making across networks)
- Mission drift under external pressure (regulatory, economic, reputational)

Structural Features:

- Convergent technology standards but divergent governance patterns
- Shared crisis response protocols
- Opt-in to inter-network ethical oversight or cultural calibration mechanisms

11.5.5 General Scaling Heuristics

Across all phases, the HAO should be guided by scaling heuristics that reflect its core philosophy:

Heuristic	Purpose
Grow by coherence, not size	Prioritize narrative, purpose, and alignment
Distribute power, not tools	Avoid central dependency on any single integration hub
License replication, not control	Encourage forking, mutation, and divergence as valid expressions of scale
Build for interop, not monoculture	Enable heterogeneous systems to coordinate through interface and protocol layers

This phased, iterative approach to scaling avoids concentrating administrative control in a single coordinating body as the network grows. Each phase adds governance capacity at the edges rather than at the center.

§11.6 Patterns of Interoperability

As Humanized Autonomous Organizations (HAOs) — the network’s coordinating framework — scale, interoperability matters more than uniformity. Traditional institutions achieve consistency through central control and standardization; HAOs achieve coherence through shared protocols, cultural schemas, and trust-based bridges across autonomous units. This final section of the implementation roadmap defines how interoperability works across four dimensions: governance, economic infrastructure, semantic coordination, and public interface integration.

11.6.1 Governance Interoperability

Governance interoperability lets different UMEs (small, self-managing venture teams), SEPs (joint ventures between teams), and distinct HAO instances recognize and trust each other’s decision-making structures without adopting the same governance forms.

Core Mechanisms:

- **Meta-Governance Protocols:** Shared templates like the Dynamic Enterprise Agreement (DEA) — a versioned operating agreement replacing fixed bylaws — serve as a common “constitutional interface,” even where specific governance logic differs across UMEs or SEPs.
- **Cross-Node Delegation:** Enables one UME to assign governance rights temporarily to another for shared initiatives or conflict mediation.
- **Multi-Sovereign Consent Models:** Allow decisions to require ratification from multiple autonomous governance bodies (e.g. a joint consent from two HAO nodes for a shared treaty).
- **Cultural Alignment Metrics:** Provide reputation-informed signals about a unit’s governance health or participation quality, feeding into trust scores.

Outcome: Autonomy is preserved, yet alignment becomes legible and actionable.

11.6.2 Economic Interoperability

As the network diversifies economically, it allows for cross-boundary financial activity while maintaining internal rules around transparency, equitable flow, and reinvestment.

Mechanisms:

- **Modular Revenue Sharing Agreements:** Allow UMEs in different legal jurisdictions or economic systems to collaborate through adaptable SEP templates.

- **Slice Conversion Interfaces:** When contribution models differ (e.g. one HAO using time-based slices, another using reputation-weighted equity), conversion interfaces allow consistent exchange and accounting.
- **Mutual Credit Bridges:** Federated mutual credit networks allow internal currencies to be converted or netted across HAO boundaries.
- **HAO Treasury Protocols:** Define how capital or resource surpluses can be invested cross-network with mutual oversight and conditional revocation.

Outcome: Money and value move across sovereign zones while remaining subject to the network’s transparency and reinvestment rules.

11.6.3 Semantic & Identity Interoperability

Scaling systems need shared language and mutual recognition across time zones, cultures, and domains, so that roles, responsibilities, and processes are intelligible to participants in different HAOs.

Components:

- **HAO Ontology Registry:** Shared vocabulary for describing roles, governance types, value flows, and lifecycle states across UMEs and SEPs.
- **Credential Translation Layers:** Allow a contributor’s reputation or credentials in one HAO (e.g., the MTU — the network’s credit-union-like financial institution) to be valid and recognized in another (e.g., the ICN — the reference cooperative business network).
- **Narrative Commons:** Story-based repositories that trace the provenance, values, and cultural logic of UMEs and SEPs in a form that’s human-readable and machine-queryable.
- **Consent and Conflict Annotations:** Allow decision records or disagreements to carry forward across contexts with embedded metadata, so that downstream actors can assess context rather than re-litigate.

Outcome: Shared vocabulary and translation layers let participants interpret each other’s roles and records without adopting identical processes.

11.6.4 Public Interface Interoperability

Interoperability with the external world (traditional markets, regulatory systems, public institutions) requires buffer systems that translate between internal and external conventions.

Approaches:

- **Federated Public Market Interfaces:** Like Contribulo, but replicated in various domains (logistics, real estate, fintech), these SEPs offer stable economic access points for investors and partners while preserving internal logic.
- **Interoperable Legal Structures:** HAOs use flexible legal wrappers (LLCs, co-ops, foundations) based on geography but align internally via the DEA. Cross-border templates reduce legal translation costs.
- **Open API Layer:** Enables integration with external software (accounting tools, supply chain software, CRMs) while enforcing network-level ethical policies through middleware.
- **Reputation/Trust Exchange:** HAOs can expose selected Value Alignment Monitoring (VAM) — ongoing checks that actions match stated principles — scores or behavioral trust histories to public partners, under zero-knowledge or tiered-access models.

Outcome: HAOs remain legible and credible to outside partners, without adopting external control mechanisms that would compromise autonomy or mission.

11.6.5 Interoperability Failure Modes and Recovery Patterns

True interoperability requires error tolerance and graceful degradation. Anticipated failure modes include:

Failure Mode	Mitigation Pattern
Semantic Drift	Periodic ontology calibration via shared narratives
Economic Imbalance	Treasury cross-subsidization with sunset clauses
Governance Asymmetry	HAO-to-HAO treaties with arbitration fallback
Identity Fragmentation	Universal credential protocols with revocation trees
Tooling Divergence	Modular open-source reference implementations

11.6.6 Emergent Pattern: The Inter-HAO Compact

As more HAOs come online, a voluntary coordination layer may emerge: a constellation of independent HAO systems that cooperate through treaties, shared ontologies, and mutual audit systems.

This is not a meta-HAO or a central authority. It consists of:

- A registry of treaty patterns and templates
- A neutral dispute resolution layer

- A cultural calibration protocol that prevents value drift
- A shared observatory for antifragility and trend sensing

The compact provides shared coordination without command authority: a federation of networks with no central governing body.

Conclusion: Interop as a Design Principle

Interoperability in HAOs functions as a design principle rather than an add-on: the model treats autonomy and connection as compatible, coordinated through cultural, technical, economic, and semantic interfaces. The HAO is intended to scale through shared purpose and protocol-based trust rather than through domination or replication of a single center.

§11.7 Formation Evidence

§11.3 describes the genesis process a Humanized Autonomous Organization (HAO) — the network’s coordinating framework — expects its first United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people — to follow: a Letter of Intent with no stated duration, then a Provisional Operations stage with an explicit three-to-six-week target window. Nothing in the deployment chapter states how long entity formation itself — the legal and financial mechanics underneath that first step — actually takes in practice. This section supplies the corpus’s only real-world timing data on that question, drawn from two source documents: a timestamped log of a single entity-formation process, and a document describing the model as it was presented to an outside pitch audience. Both describe a single instance and are presented as such, not as a general rule.

11.7.1 A Measured Formation Timeline

The corpus’s only measured entity-formation datapoint comes from a same-day log of forming provide.io as an Oregon limited liability company (LLC). The log’s stated purpose was to measure how long forming an entity could take, from no entity in existence to an account ready to receive funds. Over the course of one calendar day, the following was completed: the LLC filed with the Oregon Secretary of State, a domain registered, an Employer Identification Number (EIN) obtained from the US Internal Revenue Service (IRS), a Google Workspace account provisioned, and a business bank account opened with Mercury and approved.

Timestamped sequence, preserved as recorded:

Time	Event
Shortly after midnight	Oregon LLC filing submitted to the Secretary of State
Morning, before the Articles of Organization arrived	EIN obtained from the IRS
2:34 PM	Articles of Organization received
1:44 (the source does not specify AM or PM)	Bank account application begun
3:14 PM	Bank account application submitted
3:25 PM	Email received requesting additional information
About 5:30 PM	Additional information submitted
6:23 PM	Approval received

Two sequencing details are worth preserving alongside the timestamps. First, the EIN was obtained in the morning, before the Articles of Organization had been received that afternoon — the log records this as possible because the IRS’s registration step records a start month rather than an exact start date, so it did not require the Articles’ formal issuance timestamp to proceed. Second, the bank account application was begun at 1:44, before the 2:34 PM Articles arrived, but was not submitted until 3:14 PM, after them; the source leaves the application open across that gap rather than describing it as a continuous action.

Elapsed time, computed here from the timestamps above; the source does not state these as totals. From the Articles being received (2:34 PM) to bank approval (6:23 PM) is approximately three hours and forty-nine minutes. From the additional information being submitted (about 5:30 PM) to approval (6:23 PM) is approximately fifty-three minutes. From the initial filing (shortly after midnight) to approval (6:23 PM), the entire sequence falls within a single calendar day, well under twenty-four hours.

Privacy note. As part of opening the bank account, the applicant answered a set of identity-verification questions. Those questions asked about, in category rather than in answer: a description of the company’s business activity; a description of the specific product or service offered; external professional or social-profile links that could help verify the company or its owners; and the expected completion date of the company’s website, with whether a preview version existed. The applicant’s answers are not reproduced here, and are not otherwise summarized or paraphrased.

Relation to §11.3.3. The genesis process’s only stated duration is a design target, not a measurement: Stage 2 (Provisional Operations) is set at three to six weeks. This formation log describes an earlier and narrower slice of the process — the legal and financial mechanics that would sit at or before Stage 1 (Intent to Form), which itself carries no stated duration in §11.3.3 — and shows that slice completing within a single calendar day in this one instance. The two figures answer different questions and are not directly comparable: one is a multi-week target for an operations stage that follows formation, the other is a same-day measurement of formation itself.

11.7.2 A Three-Stage Model Presented to an External Audience

A separate document, dated by its own front matter to 2024-07-07, records how the model was described to an audience outside the network: a submission to Business Impact NW's IMPACT Pitch 2024 competition. Because it was written to win a pitch rather than to specify a design, its claims are read here as what an applicant told a funding audience, not as a position the corpus otherwise adopts. This is the corpus's only record of the model presented to an external audience in this form.

The document lays out a three-stage rollout with explicit calendar months attached — the only place in either source consulted for this section that ties a phased rollout to elapsed time rather than to headcount or milestone:

- **Think Tank (Months 0–6):** concept validation, model development, and educational content creation.
- **Incubator (Months 7–12):** supporting the formation of collaborative businesses, tracking adoption, customization, and early sustainability.
- **Integration Services Provider (Months 13–18):** measuring and scaling impact, including comparison against industry benchmarks and tracking of local economic effects.

Elsewhere in the same document, the same three stages are named “Think Tank,” “Incubation,” and “Integration,” without the month ranges attached — one document uses two slightly different namings for the same three-stage structure.

The document describes a single planned application meant to combine several functions into one flow: helping people find collaborators, providing education, running identity-verification checks, starting a new business, and managing an existing one — framed as comparable in ease to using an online banking application. That identity-verification step describes, at the level of a product concept, the same category of task that §11.7.1's formation log documents as a manual, multi-hour exchange with a bank. The two records share a provide.io context but do not reference each other directly; the connection is drawn here as a plausible link between a stated product vision and a lived process, not as a fact either source states.

The pitch also describes a micro-learning feature grounded in behavioral science, aimed at skills including emotional intelligence, conflict resolution, leadership, and business management, delivered through game-like elements.

The document states an explicit rationale for using LLCs rather than cooperative or corporate structures for individual ventures: formation speed and legal portability across states. It is not the only place the corpus gives a reason for an LLC wrapper: docs/08-legal/01-modular-legal-forms.md lists “easy formation, broad use, member-managed options” as LLC strengths against limits including “not inherently cooperative, lacks embedded mission.” What is specific to this document is that the rationale is offered by an applicant to an outside audience rather than stated as design. §11.2 lists “Flexible Benefit LLC or Platform Cooperative” as alternative legal wrappers without stating a reason to choose between them.

On geographic focus, the document states an initial scope of proof-of-concept efforts that are remote or located in Portland, Oregon, continuing until judged viable by an unspecified vote. No further siting rationale is given, and the document does not define what the viability vote consists of.

11.7.3 A Single Instance, Not a Base Rate

Both records in this section describe one instance each: one measured formation process for one entity, one pitch document written for one audience at one point in time. Neither is a repeated or independently corroborated measurement. The corpus states durations elsewhere — §11.3's three-to-six-week stage target, and feasibility and pilot windows in chapters 04, 06, and 13 — but those are design targets and planning assumptions rather than measurements, so none of them is a comparable against which either record here can be checked.

§11.3's genesis process states a design target for one stage of onboarding without measuring it. §11.7.1 measures one instance of an earlier stage without stating a target for it. The two are complementary rather than substitutable: a target window states an expected duration for a designed process stage; a single timestamped log shows that a related, earlier legal and financial process completed within one day at least once. Neither can stand in for the other, and neither should be read as establishing what is typical. One same-day formation demonstrates that fast formation is possible; it does not establish an average, a rate, or a planning assumption the network can rely on. A larger, repeated sample of entity formations would be needed to state a typical duration for the legal and financial layer with any confidence.

The pitch document's three-stage rollout (§11.7.2) sits on the same footing: it is a stated intention presented to a funding audience, not a measured outcome. Nothing in either source establishes that the Months 0–6 / 7–12 / 13–18 schedule was met.

12.0 • ICN

Introduction

The Integrated Cooperative Network (ICN) — the reference cooperative business network — is an implementation of the Humanized Autonomous Organization (HAO) — the network’s coordinating framework — model, designed to distribute value creation, decision-making authority, and revenue flow across diverse, semi-autonomous entities. Where the HAO specifies the architecture, the ICN provides an implementable pattern that combines polycentric governance, trickle-out economics, and defined governance safeguards into a socio-technical system.

At its core, the ICN is built around **United Micro Enterprises (UMEs)**: small, self-governing organizational units that concentrate value generation at the edges of the network. These UMEs are interlinked through **Strategic Enterprise Partnerships (SEPs)**: collaborations formed to pursue shared goals, pool risk, and combine capacity. The network’s coordination and infrastructure logic is maintained by the **Humanized Autonomous Organization (HAO)** layer, which distributes protocols, governance tools, and support systems across the ecosystem rather than centralizing control.

Economically, the ICN reverses the typical direction of capital flows by implementing a **trickle-out architecture**: investment flows directly to productive nodes (UMEs). Separately, diminishing contributions return a declining portion of their revenue to shared network infrastructure over time. This is intended to support early-stage UMEs while limiting long-term dependence on, and concentration of, network-level economic power.

The ICN also introduces the Micro Enterprise Ecosystem (MEE) — the network’s protected internal economy: a semi-permeable boundary that buffers value creators from external market volatility while enabling selective, controlled interfacing. Public Market Interfaces (PMIs) — a buffer company between the network and outside investors — such as the *Contribulo* example, translate between the ICN’s internal logic and conventional financial expectations.

This section details the ICN’s structure, governance protocols, economic flows, technological systems, and implementation dynamics. The ICN is specified in enough detail to be implemented: the architecture can be forked, adapted, and recontextualized by communities implementing this model.

The ICN implements the HAO model through coordinated autonomy, defined technology governance, and distributed ownership.

§12.1 Structural Features: UMEs, SEPs, MEE, and HAO

The Integrated Cooperative Network (ICN) — the reference cooperative business network — is organized as a fractal, polycentric structure of decentralized agency, modular collaboration, and adaptive coordination, rather than a hierarchical corporate model. This section introduces the four primary structural components: **United Micro Enterprises (UMEs)** — small, self-managing venture teams ($\leq \sim 15$ people); **Strategic Enterprise Partnerships (SEPs)** — joint ventures between teams; the **Micro Enterprise Ecosystem (MEE)** — the network’s protected internal economy; and the overarching **Humanized Autonomous Organization (HAO)** — the network’s coordinating framework.

12.1.1 United Micro Enterprises (UMEs)

UMEs are the fundamental, value-generating units of the ICN: semi-autonomous, self-governing economic entities, each with its own governance structure, operational processes, and internal culture. The design principle underpinning UMEs is **bounded autonomy**: they operate independently within a framework of shared principles, infrastructure, and accountability mechanisms.

Key characteristics of UMEs include:

- **Scale Constraint**: As an ICN design heuristic, UMEs are intentionally small — typically no more than 8–15 active contributors — to support trust, agility, and internal coherence. The range is a corpus design choice, not an empirical finding.
- **Embedded Governance**: Each UME selects its preferred internal governance mode (e.g., consent-based sociocracy, rotating stewardship, or role-based decision authority), as long as it remains compliant with the ICN’s Dynamic Enterprise Agreement (DEA) — a versioned operating agreement replacing fixed bylaws.
- **Autonomous Capitalization**: While initial investments may be seeded via the HAO, UMEs maintain their own capital structures, accounting systems, and reserve funds. They may attract local investments, issue internal equity to members, or participate in cooperative lending systems.
- **Lifecycle Independence**: UMEs follow a defined lifecycle: **Genesis, Operational Maturity, Replication or Aggregation, Degradation, and Dissolution**.

12.1.2 Strategic Enterprise Partnerships (SEPs)

SEPs are contractual, purpose-driven collaborations between two or more UMEs. While UMEs focus on localized value creation, SEPs enable horizontal scaling through cooperative specialization.

Key features of SEPs:

- **Joint Mission Alignment**: SEPs are initiated around a specific opportunity or problem domain. They may produce a shared product, infrastructure layer, service bundle, or joint market entry strategy.
- **Custom Governance Layer**: A lightweight, dynamic governance mechanism is used to manage the SEP, distinct from the internal governance of participating UMEs. This includes defined decision rights, dispute resolution processes, and exit protocols.

- **Revenue Sharing Agreements:** SEP economics are governed by a signed, versioned **SEP Agreement**, which defines contribution weights, shared costs, IP handling, and reinvestment terms.
- **Reinvestment Pools:** Many SEPs create an internal reinvestment pool for collective innovation or resilience-building, akin to a micro-venture fund for the partnership.

SEPs resemble mycorrhizal networks: inter-organizational linkages that let entities exchange value, signals, and surplus without merger or acquisition.

12.1.3 Micro Enterprise Ecosystem (MEE)

The MEE is the protected socio-economic environment within which UMEs and SEPs operate, forming a semipermeable boundary layer between the ICN and external economic systems.

Key structural roles of the MEE:

- **Selective Permeability:** Access to the ecosystem is regulated through **Enterprise Integration Assessments (EIA)**. This process checks that new participants or external interfaces meet the ICN’s ethical and operational alignment standards.
- **Market Buffering:** The MEE buffers internal actors from volatile market pressures, particularly when interfacing with legacy financial institutions, venture capital, or concentrated market power.
- **Internal Commons Access:** UMEs within the MEE can access shared infrastructure (open software libraries, design assets, training modules, and collective legal services) through **Commons Protocols** maintained by the HAO.
- **Cultural Integrity Zone:** Through tools such as **Enterprise Culture Cultivation (ECC)** and **Value Alignment Monitoring (VAM)** — ongoing checks that actions match stated principles — the MEE supports internal coherence while allowing local variation.

The MEE functions as a voluntary, values-based perimeter that supports collective resilience and autonomy without central control.

12.1.4 Humanized Autonomous Organization (HAO)

The HAO’s role in the ICN is to maintain infrastructure, protocols, and cohesion across the network, without functioning as a command structure.

Core functions of the HAO in the ICN context:

- **Protocol Maintenance:** The HAO authors and version-controls core governance frameworks such as the DEA, economic protocols, and SEP templates. These form the ICN’s operating system.
- **Infrastructure Stewardship:** It maintains and evolves the **Distributed Ledger Infrastructure (DLI)**, **Collaborative Intelligence Network (CIN)**, and other platform services.
- **Capital Flow Orchestration:** The HAO initiates initial investments into UMEs, manages pooled reserves, and redistributes excess capital through multi-tiered distribution schedules (monthly operational, quarterly reconciliations, annual strategy-driven redistributions).
- **Conflict Resolution and Arbitration:** It offers optional third-party mediation services, backed by trained practitioners, to handle ethical or operational disputes within the ICN.
- **Evolution and Sensing:** The HAO conducts Enterprise Integration Assessments (EIA), tracks cultural drift, and initiates participatory reconfiguration processes intended to support long-term alignment.

The percentage of revenue flowing to the HAO declines over time, per the **Diminishing Contribution Mechanism**, as UMEs repay startup investments and reach operational maturity — typically from 30–40% (early stage) to 10–15% (mature stage).

The HAO coordinates the network’s infrastructure, capital flows, and governance processes without direct command authority over UMEs.

Summary Table: ICN Structural Layers

Layer	Function	Scope	Autonomy	Governance
UME	Value creation, small-team autonomy	Local	High	Internal + DEA compliant
SEP	Joint ventures between UMEs	Cross-UME	Medium	Joint SEP Agreement
MEE	Protected economic and cultural zone	Network-wide	Regulated	Maintained by HAO
HAO	Infrastructure and meta-governance	System-wide	N/A	Participatory, protocol-driven

This structural configuration allows the ICN to scale through network multiplication — adding UMEs and SEPs — rather than through aggregation into larger centralized units.

§12.2 Governance Framework and DEA Protocols

Governance in the Integrated Cooperative Network (ICN) — the reference cooperative business network — is designed to be adaptive, distributed, and evolvable. Decision-making authority is distributed across nested levels of autonomy — UMEs (small, self-managing venture teams, $\leq \sim 15$ people), SEPs (joint ventures between teams), and cross-network assemblies — rather than centralized in the Humanized Autonomous Organization (HAO), the network’s coordinating framework. This design is codified in the **Adaptive Governance Framework (AGF)** — the network’s layered governance system — and instantiated in a version-controlled governing document called the **Dynamic Enterprise Agreement (DEA)** — a versioned operating agreement replacing fixed bylaws.

12.2.1 The Adaptive Governance Framework (AGF)

Ostrom (2010) describes polycentric systems as multiple formally independent decision centers capable of coherent coordination. The AGF applies that premise, alongside sociocratic consent models and complex adaptive systems thinking, as a corpus design for balancing decentralized autonomy with network-wide coherence.

Core design features:

- **Polycentric Authority:** Multiple centers of decision-making power exist concurrently: UMEs govern themselves internally, SEPs negotiate terms jointly, and the HAO facilitates governance of shared infrastructure and protocol evolution. No layer has unilateral control.
- **Subsidiarity Principle:** Decisions are made at the lowest competent level. Issues affecting only a single UME are resolved internally; inter-UME matters are handled via SEP-specific agreements; systemic issues are escalated to HAO-coordinated assemblies or protocol changes.
- **Consent-Based Decision Making:** Rather than majority vote or consensus, most ICN-wide decisions use **consent**: a decision is valid unless there is a reasoned objection that cannot be resolved, allowing dissent to be raised without blocking all progress.
- **Nested Governance:** Governance is layered, with different domains of authority:
 - *Constitutional Layer:* HAO and DEA protocols
 - *Strategic Layer:* SEP-level agreements and system-wide initiatives
 - *Operational Layer:* UME-level decisions, project execution, and internal role distribution
- **Feedback Integration:** The AGF incorporates feedback loops — participatory sensing, alignment audits, and evolutionary protocols — that let the network revise its own rules through deliberation and versioning.

The AGF is designed so that power remains contextual, responsive, and non-accumulative, supporting autonomy at the edge without sacrificing systemic stability.

12.2.2 The Dynamic Enterprise Agreement (DEA)

The DEA is the ICN’s version-controlled governing document. Unlike static bylaws or charters, it evolves alongside the network through transparent, traceable versioning.

Key attributes of the DEA:

- **Versioned Structure:** Each revision of the DEA is tagged, timestamped, and includes a changelog, rationale, and signatory verification, creating a record that supports forensic analysis, dispute resolution, and institutional memory.
- **Modular Sections:** The DEA is composed of interoperable modules:
 - *Core Values and Purpose*
 - *Governance Mechanisms and Authority Maps*
 - *Economic Protocols and Contribution Formulas*
 - *Membership Definitions and Participation Rights*
 - *Conflict Resolution Procedures*
 - *Integration and Exit Protocols*
- **Participant Scope:** All UMEs and SEPs are required to maintain DEA compliance to remain within the ICN’s Micro Enterprise Ecosystem (MEE) — the network’s protected internal economy. They are not bound by the DEA in their internal governance, only in how they interface with the network.
- **Evolution Protocols:** Changes to the DEA follow a structured process:
 1. **Proposal Initiation:** Any UME or member may propose a change
 2. **Deliberation Phase:** Proposals are evaluated by a randomly selected or rotating **Deliberation Cell**
 3. **Consent Phase:** Broader network consent is sought through multi-channel engagement
 4. **Version Lock-In:** If ratified, the new DEA version is cryptographically signed and distributed

- **Access and Transparency:** The DEA is public and machine-readable, published in formats (e.g., Markdown, JSON-LD, IPFS-based) parsed by humans, legal systems, and smart contracts alike.

12.2.3 Trust, Roles, and Representation

The ICN replaces board or shareholder governance with trust-weighted participation and dynamic role structures:

- **Roles, Not Titles:** Operational authority is allocated to clearly defined roles, not static titles. Roles are bound to responsibilities, scope of authority, and revocation mechanisms.
- **Trust as Credentialing:** Participation in sensitive or high-impact decisions (e.g., DEA evolution, SEP arbitration) may require demonstrated trust history, based on:
 - Length and quality of contributions
 - Prior role performance
 - Peer endorsements
 - Audit logs and value alignment metrics
- **Representation Systems:**
 - *Rotational Cells:* Temporary working groups selected randomly or via rotating nomination to ensure diversity of voice.
 - *Domain Delegates:* Appointed or elected based on relevant expertise, time commitment, and trust score.
 - *Deep Democracy Channels:* Anonymous polling, deliberative forums, and “sensing” tools capture perspectives beyond formal roles.

This system ties legitimacy to demonstrated trust and participation rather than formal title.

12.2.4 Conflict Resolution and Grievance Handling

All ICN actors agree to a network-wide **Graduated Conflict Resolution Protocol (GCRP)**:

1. **Direct Engagement:** Attempt resolution between parties
2. **Facilitated Dialogue:** Involve a neutral party or trained facilitator
3. **Mediation Circle:** Invite a temporary governance microcell to assist
4. **Network Arbitration:** Submit case to a standing deliberation or ethics cell
5. **Exit or Partition:** In rare, unresolved cases, initiate opt-out or structured exit processes per DEA

This approach emphasizes restoration over punishment, aiming to keep governance relationship-centric and repairable.

12.2.5 Governance Summary Table

Mechanism	Scope	Method	Evolvability	Example
AGF	System-wide	Polycentric, nested, subsidiarity	Reflexive and modular	Decision routing logic
DEA	Network-wide compliance	Version-controlled, modular charter	High – structured evolution path	Governance contract history
UME Governance	Local only	Chosen by UME (e.g., consent, roles)	Internal to UME	Sociocratic circles
SEP Agreement	Cross-UME project-specific	Negotiated contractual terms	Typically project-limited	SEP IP agreement
Conflict Handling	Multi-layered	Graduated resolution protocol	Evolvable via DEA	Grievance response escalation

By embedding evolvability, distributed decision-making, and accountability at every layer, the ICN’s governance model differs from static institutional bylaws and from code-only DAO governance: rules can be revised through deliberation and versioning rather than fixed at founding or set entirely by pre-written code.

§12.3 Economic Model: Trickle-Out Capital Allocation, Diminishing Contributions, and Equity

The Integrated Cooperative Network (ICN) — the reference cooperative business network — uses a **trickle-out economic model** to allocate investment outward toward UMEs (small, self-managing venture teams, $\leq \sim 15$ people) and SEPs (joint ventures between teams) first. Separately, diminishing contributions return a declining share of revenue to shared coordination infrastructure as those units mature. Together, these mechanisms are intended to reward early contribution and align incentives across the network’s layers: UMEs, SEPs, and the Humanized Autonomous Organization (HAO) — the network’s coordinating framework.

Economic flows are designed to track actual contribution and maturity over time rather than fixed cost-center accounting. This section outlines the ICN’s investment logic, revenue allocation framework, reinvestment mechanisms, and equity logic.

12.3.1 Capital Entry and Investment Allocation

All external and internal capital first enters the ICN through the HAO layer, which routes capital rather than holding economic power directly.

Capital Entry Points:

- Initial donor or investor funding
- Revenue from Public Market Interfaces (PMIs) — buffer companies between the network and outside investors
- Reinvestments from mature UMEs or SEPs
- Platform revenues (e.g., software, shared infrastructure licensing)

Directed Investment Mechanism: Upon entry, funds are redirected to value-producing UMEs and collaborative SEPs based on:

- **Strategic Alignment:** Alignment with the ICN’s long-term goals
- **Maturity Phase:** Early-stage units receive higher support
- **Capability Assessment:** The team’s execution readiness
- **Opportunity Index:** Market or mission impact potential
- **Readiness Score:** Operational capacity and infrastructure in place

This mechanism inverts the conventional top-down direction of capital flow, directing investment to UMEs and SEPs before it reaches central coordination.

12.3.2 Diminishing Contribution and Revenue Allocation Framework

Diminishing contributions from revenue generated by UMEs return to shared HAO infrastructure through a dynamic contribution model. The percentage contributed declines over time as UMEs mature and repay their startup investment.

Lifecycle-Based Distribution:

UME Stage	HAO Allocation	UME Retention	Notes
Early Stage	30–40%	60–70%	Heavy infrastructure support; startup phase
Mid Stage	20–25%	75–80%	Stable operations, repayment underway
Mature Stage	10–15%	85–90%	Fully sovereign, minimal central reliance

The revenue sent to the HAO is allocated to systemic reinvestment rather than distributed as profit.

Distribution Frequency:

- **Monthly:** Operational distributions
- **Quarterly:** Performance-based adjustments
- **Annually:** Strategic redistributions

These schedules maintain predictable cash flow while enabling adaptive network-wide financial planning.

12.3.3 SEP-Level Economic Agreements

When two or more UMEs enter a SEP, they establish a joint revenue-sharing agreement, negotiated in advance.

Key SEP Economic Elements:

- **Contribution Valuation:** Time, assets, IP, and risk-weighted effort
- **Revenue Distribution Schedule:** May be fixed, milestone-based, or percentage-based
- **Reinvestment Clause:** Optional % of revenue locked into a SEP pool
- **HAO Share:** Lower than direct UME contributions; typically 5–10% to maintain shared infrastructure

This design supports collaborative economics without requiring mergers or acquisitions, preserving autonomy while scaling capacity.

12.3.4 Reinvestment and Mutual Capital Pools

A percentage of all revenue distributed to the HAO is directed into network reinvestment mechanisms rather than distributed as profit outside the network.

HAO Reinvestment Purposes:

- Incubate new UMEs or SEPs
- Provide support for struggling or high-risk units
- Fund shared infrastructure (tech, legal, design)
- Expand into underserved bioregions

- Provide mutual aid and emergency relief

Mature UME Reinvestment: Mature UMEs may opt into voluntary reinvestment, such as:

- **Mentorship Funding**
- **Rotating Innovation Funds**
- **UME-to-UME Credit Lines**
- **Contribution to shared Commons pools**

This system recirculates value within the network and reduces reliance on external capital.

12.3.5 Equity Design: Distributed Ownership and Alignment

Unlike corporate equity models that concentrate control in founders or investors, the ICN uses a dynamic, hybrid equity structure based on contribution-weighting and long-term alignment.

Key Principles:

- **Slices Not Shares:** The ICN may calculate equity using a *Slicing Pie*-style model (Moyer, 2012): each member earns “slices” based on at-risk time, money, IP, or other contributions. Treating network value as a contribution category is a corpus extension.
- **Multipliers for Risk:** Contributions are weighted (e.g., 2x for time, 4x for cash) based on risk and post-tax scarcity.
- **Vesting Mechanism:** Equity accrues gradually and vests through continued participation and alignment.
- **Diminishing HAO Stake:** The HAO may retain a non-voting equity slice initially, which gradually dissolves as investment is repaid.

Equity Accrual Formula:

$$\text{Slices} = \text{Contribution} \times \text{Risk Multiplier}$$

$$\% \text{ Equity} = (\text{Your Slices} / \text{Total Slices}) \times 100$$

This formula ties equity to measured contribution rather than to initial capital or founder status.

12.3.6 Member Compensation Structure

Members of UMEs are compensated through a four-tier model:

1. **Base Compensation** – Regular income for stability (hourly, salary, or retainer)
2. **Performance-Based Distribution** – Monthly or quarterly bonuses linked to UME health
3. **Profit Sharing** – Annual or milestone-based, tied to collective outcomes
4. **Equity Accrual** – Long-term ownership stake in the UME and/or SEP

This model balances:

- Immediate financial security
- Medium-term performance incentives
- Long-term wealth-building
- Ethical, transparent reward mechanisms

Compensation is structured to link individual reward to the UME’s long-term value rather than short-term payout.

Summary: Economic Philosophy Comparison

Model	Traditional Corp	DAO Model	ICN Trickle-Out Model
Capital Flow	Top-down	Treasury-managed	Edge-directed (UMEs first)
Revenue Use	Profit extraction	Token inflation or burn	Reinvestment and redistribution
Equity	Founder/investor concentrated	Token-based, volatile	Dynamic, contribution-weighted
Compensation	Salary + bonus	Token airdrops or bounties	Layered, contribution-weighted, long-term
Scaling	Capital accumulation	Protocol forking	Network multiplication via UMEs

The ICN’s economic model links capital, labor, coordination, and ownership through incentive design embedded directly in the organizational architecture, structured as an alternative to platform-capitalism and token-based DAO economic models.

§12.4 Technical Stack and Platform Design

The technical infrastructure of the ICN — the reference cooperative business network — is a modular, interoperable stack designed to support distributed autonomy, coordinated collaboration, and ethical computation, rather than a monolithic platform. The architecture emphasizes human oversight over algorithmic determinism, supporting trust, transparency, and adaptability across **United Micro Enterprises (UMEs)** — small, self-managing venture teams ($\leq \sim 15$ people) — **Strategic Enterprise Partnerships (SEPs)** — joint ventures between teams — and the **Humanized Autonomous Organization (HAO)** — the network's coordinating framework.

This section outlines the stack's foundational components, their integration logic, and the platform-level services provided to the Micro Enterprise Ecosystem (MEE) — the network's protected internal economy.

12.4.1 Design Principles of the ICN Stack

The ICN's technical architecture is governed by five design axioms:

1. **Distributed First:** No single point of failure or control. Infrastructure must support federation, redundancy, and localization.
 2. **Composable and Extensible:** All systems are modular, with exposed APIs, schemas, and plugin support for local customization and rapid evolution.
 3. **Human-in-the-Loop by Default:** Critical operations (e.g., value alignment audits, governance decisions) require interpretability and human discretion rather than fully automated execution.
 4. **Selective Transparency:** Transparency internally, with selectively permeable external interfaces. Privacy and consent are foundational.
 5. **Ethical Traceability:** Every major decision, transaction, or governance update is traceable with justification metadata, intended to support accountability without functioning as surveillance.
-

12.4.2 Core Infrastructure Components

The following components form the core of the ICN's infrastructure layer, coordinated by the HAO and accessible to all UMEs and SEPs.

a. Distributed Ledger Infrastructure (DLI)

- **Function:** Immutable, versioned record-keeping for governance changes, financial flows, equity accrual, and contribution histories.
 - **Stack:** Private permissioned ledger (e.g., Hyperledger Fabric), integrated with public-facing gateways for select audits or public PMIs (buffer companies between the network and outside investors).
 - **Use Cases:**
 - DEA (a versioned operating agreement replacing fixed bylaws) version control and signing
 - Revenue sharing verifiability
 - Cross-UME SEP contract records
-

b. Collaborative Intelligence Network (CIN)

- **Function:** Augmented decision-making platform combining human insight with machine intelligence.
 - **Stack:** Distributed knowledge graphs, AI inference engines, consent-based machine learning feedback loops.
 - **Use Cases:**
 - Conflict anticipation and scenario modeling
 - Dynamic governance simulations
 - Skill-matching and role suggestion within UMEs/SEPs
-

c. Value Alignment Monitoring (VAM) — ongoing checks that actions match stated principles

- **Function:** Continuous, passive tracking of alignment between operational decisions and ICN values (e.g., ETHICAL framework).
 - **Stack:** Declarative intent statements, sentiment analysis, operational data tagging.
 - **Use Cases:**
 - Drift detection and early warnings
 - Cultural cohesion audits
 - DAO-style proposal moderation
-

d. Enterprise Culture Cultivation (ECC)

- **Function:** Tools and frameworks for shaping shared language, values, and collaboration styles across UMEs.

- **Stack:** Ritual design templates, onboarding flows, micro-learning portals, story-driven culture tools.
 - **Use Cases:**
 - UME identity development
 - MEE-wide narrative coherence
 - Customizable local practice toolkits
-

e. Trust Signaling and Verification Layer

- **Function:** Infrastructure to evaluate and signal multi-dimensional trust (operational, social, financial) across the network.
 - **Stack:** Layered reputation graphs, verifiable credentials, peer endorsement protocols.
 - **Use Cases:**
 - SEP partner selection
 - Role delegation and access control
 - Financial risk balancing in the style of the MTU (the network's credit-union-like financial institution)
-

12.4.3 Platform-Level Services

The HAO maintains shared services accessible to all UMEs and SEPs via authenticated, permissioned access:

- **Commons Registry:** Shared assets (e.g., legal templates, design systems, software libraries)
 - **Governance Engine:** DEA revision interface, SEP contract templates, deliberation protocols
 - **Equity Ledger:** Tracks and updates dynamic ownership via slices across UMEs and SEPs
 - **Conflict Resolution Hub:** Routing and mediation tools for resolving internal or cross-UME disputes
 - **Onboarding + Learning Systems:** Personalized onboarding flows, capability assessments, and self-directed development tracks
 - **Metrics Dashboard:** Multi-capital performance indicators, participatory audits, and alignment signals
 - **Public Market Interface Gateways:** Trusted APIs for PMIs to access authorized, aggregate ICN data (financials, growth metrics, etc.)
-

12.4.4 Integration with External Systems

The ICN stack supports selective interoperability with public blockchains, legacy enterprise software, and government regulatory systems, where aligned with network goals.

Examples:

- **CoopCycle-style license integration** for logistics SEPs
- **eIDAS or SSI integration** for identity and legal compliance
- **OAuth/OIDC bridges** to allow federated login across UMEs
- **Accounting APIs** to map ICN ledger outputs into QuickBooks or other local systems

Where alignment is not possible (e.g., platforms built on extensive user data collection, or financial APIs with incompatible fee structures), economic firewalls are enforced through MEE boundary logic and policy-based data contracts.


12.4.5 Reference Implementation: provide.io as a Coordinating Engine

The technical design and integration architecture of the ICN can be partially embodied by provide.io, a hybrid for-profit/for-purpose infrastructure and systems integrator, which can serve as:

- **Genesis Engine:** Bootstrapping the first cluster of UMEs and SEPs
- **Protocol Steward:** Maintaining reference implementations of the DEA, DLI, and CIN modules
- **Developer Hub:** Publishing SDKs, CLI tools, APIs, and validators for contributors and edge deployers
- **Multi-Tenant Hosting Provider:** Running testnets, staging environments, or sovereign deployments of the ICN stack

While not essential to the ICN itself, provide.io demonstrates how reference implementations can accelerate deployment while remaining interoperable with alternative forks or localized adaptations.

Summary Diagram: ICN Technical Stack



The ICN’s technical platform is a modular infrastructure rather than a single application or centralized protocol: interoperable services supporting cooperation, coordination, and governance across the network, with human oversight built into critical decision points (§12.4.1).

§12.5 Public Market Interfaces (PMIs)

Bridging Regenerative Economies with Traditional Markets

The Integrated Cooperative Network (ICN) — the reference cooperative business network — is structured to limit exposure to public-market demands for short-term returns that conflict with reinvestment priorities. To achieve strategic goals such as capital acquisition, visibility, infrastructure scaling, or external supply chain integration, the ICN sometimes interacts with public markets through **Public Market Interfaces (PMIs)** — buffer companies between the network and outside investors — intermediary structures that translate between ICN principles and public-market practice without compromising network integrity.

This section outlines the purpose, structure, and operational logic of PMIs, including the example of **Contribulo**, an illustrative implementation model.

12.5.1 The Function of Public Market Interfaces

PMIs function as semi-permeable membranes between the ICN’s Micro Enterprise Ecosystem (MEE) — the network’s protected internal economy — and public markets. Their core functions include:

- **Translation of Value Logic:** Converting ICN-based metrics (such as long-term trust, equitable participation, or ecological impact) into formats legible to external investors, partners, or regulatory bodies.
- **Capital Buffering:** Attracting and managing external capital without exposing UMEs (small, self-managing venture teams) or SEPs (joint ventures between teams) directly to pressure for short-term returns or to hostile takeovers.
- **Legal and Financial Insulation:** Creating firewalls between network-owned entities and public shareholders, using hybrid ownership structures and firewall clauses.
- **Reputation and Signaling:** Representing the ICN in the language expected by public audiences (annual reports, disclosures, impact KPIs) while maintaining alignment with the network’s stated values.
- **Strategic Integration:** Managing partnerships with third-party logistics, fintech platforms, B2B distributors, and regulatory institutions.

12.5.2 Structural Patterns of PMIs

PMI structures vary by purpose, jurisdiction, and context. The ICN may instantiate multiple interfaces. Common patterns include:

Type	Description	Use Case
Hybrid LLC	Traditional entity with custom operating agreement	Contribulo model
Joint Venture Entity	Co-owned by ICN and legacy actors	Logistics SEP with CoopCycle
Steward-Ownership Company	Majority control retained by mission-aligned stakeholders	Long-term infrastructure management
Open IP Licensing Body	Issues usage rights to external actors under specific conditions	Software, protocol, or logistics deployment
Public Benefit Corporation (PBC)	Legal form recognizing social/environmental objectives	External-facing advocacy and partnerships

Each PMI is governed by a **Purpose-Specific Operating Agreement (PSOA)**, which embeds:

- Value Alignment Clauses
- Return Expectations and Limits

- Dissolution/Exit Conditions
 - HAO (the network's coordinating framework) Oversight Provisions
 - Equity and Voting Protections
-

12.5.3 Example: Contribulo as a Prototype PMI

Contribulo (structured in the model as a hybrid LLC) serves as a prototype for how the ICN might interface with public markets without diluting its principles.

Illustrative Design:

- **Ownership Structure:**
 - Allocation between network stewardship and public investment remains pending an owner decision
- **Role:**
 - Manages PMIs related to professional services, contributor platforms, and revenue-generating interfaces
 - Provides liquidity access to contributors without exposing internal governance
- **Control Mechanisms:**
 - Governance veto rights retained by HAO-appointed trustees
 - Caps on dividend distributions to maintain reinvestment orientation
 - Mission lock clauses and legal obligations to ICN charter
- **Functionality:**
 - Operates a freelance or service marketplace with ICN-derived ethics and logic
 - Accepts capital and operates as a public-facing representative of the ICN in traditional markets, while buffering the MEE

❖ *Note: Contribulo is an illustrative pattern rather than a canonical requirement. Other PMIs may have entirely different compositions based on regulatory, cultural, or functional context.*

12.5.4 Economic Flow and Containment

PMIs function as economic airlocks between market capital and ICN economic flow:

- **Inbound Capital:**
 - External investment → PMI → Directed deployment to targeted UMEs or SEPs
 - Always governed by deployment contracts with pre-defined reinvestment expectations
 - **Outbound Returns:**
 - Profits accrued in the PMI can be distributed, with dividends capped and mission-aligned
 - Excess returns are redirected to HAO-controlled reinvestment pools
 - **Isolation Layer:**
 - If a PMI becomes misaligned, its link to the ICN can be severed, limiting disruption to the rest of the network
 - Firewallled contracts are designed to prevent leakage of ICN decision rights, equity, or assets
-

12.5.5 Governance and Accountability Mechanisms

PMIs are required to follow an accountability stack:

- **Internal ICN Oversight:**
 - Regular audit and ethics compliance reports submitted to the HAO
 - Embedded governance representatives from UMEs, SEPs, and members
- **Transparent Reporting:**
 - Impact metrics, return logic, and governance activities made publicly available
 - Metrics track against ETHICAL and PARTS frameworks (e.g., % local reinvestment, % equity held by contributors)

- **Revocability:**
 - The ICN retains revocation or forking rights for any PMI that no longer meets governance or ethical criteria

12.5.6 PMI Risks and Design Safeguards

While PMIs provide operational utility, they also introduce risks:

Risk	Mitigation
Value Drift	Regular VAM (ongoing checks that actions match stated principles) audits; enforceable alignment clauses
Capture by Investors	Non-voting equity, mission locks, veto rights
Reputation Contamination	Firewalled branding, public alignment dashboards
Overdependence	Diversification of PMIs; fallback protocols within MEE

Summary: The PMI Pattern

PMI Function	Legacy Analogue	ICN Difference
Capital Access	IPO, VC, Token Sale	Buffered entry via mission-aligned entity
External Engagement	Public Corp	Values-locked operating agreement
Legal Representation	ParentCo	Legally constrained role with network subordination
Revenue Translation	Platform Corp	Proceeds allocated outward through trickle-out logic
Brand Protection	PR Arm	Selectively permeable membrane + value firewalls

PMIs like Contribulo let the ICN interact with external economic systems without integrating into them, providing a structure for engaging public markets while limiting the risk of principle dilution or organizational capture. Properly designed, a PMI translates the ICN's internal logic into forms external parties can engage with, without altering the ICN's internal governance.

13.0 • Member Trust Union

Executive Summary

The **Member Trust Union (MTU)** — the network’s credit-union-like financial institution — is a domain-specific implementation of the **Humanized Autonomous Organization (HAO)** — the network’s coordinating framework. It manages financial coordination for the **Integrated Cooperative Network (ICN)** — the reference cooperative business network. Rather than credit scores, profit-maximization models, or centralized control, the MTU relies on multi-level trust verification, surplus recirculation, cultural-alignment checks, and polycentric governance.

This section describes how the MTU applies HAO principles to build financial infrastructure designed for scale and interoperability with external systems. It allocates capital, distributes risk through a mutual credit system, and translates between decentralized value creation and conventional financial systems.

13.1 – Multi-Level Trust Architecture

Establishes a nested, polycentric structure: individuals → Local MTUs → regional federations → global MTU layer. Trust verification and progressive stewardship serve as the basis for financial access and participation, replacing abstract scoring with social capital.

13.2 – Financial Product Design

Introduces purpose-aligned tools such as peer-to-peer lending, collective asset financing, timebanks, and shared equity pools. These are embedded in trust relationships and community ownership rather than market pricing.

13.3 – Governance and Risk Distribution

Describes the MTU’s polycentric, consent-based governance model, including dynamic liquidity protocols, pooled reserves, adaptive risk weighting, and mutual aid systems. Governance authority is distributed, risk is localized and buffered, and accountability mechanisms are documented and auditable by members.

13.4 – Technology and Physical-Digital Integration

Details the MTU’s hospitality-centered physical spaces and their integration with APIs, ledgers, identity systems (DIDs), and compliant fintech tooling. Digital infrastructure supplements in-person interactions rather than replacing them.

13.5 – Member Experience and Compensation Systems

Outlines how members receive base livelihood compensation, performance-aligned distributions, and equity accrual, forming a long-term economic relationship with their community and the ICN. Lifecycle tracking of trust status replaces role-based financial gatekeeping.

13.6 – Integration with ICN and External Markets

Covers internal integration with the ICN’s governance, coordination among UMEs (small, self-managing venture teams, each $\leq \sim 15$ people), and revenue sharing among SEPs (joint ventures between teams). Externally, MTUs interface with financial institutions through **Public Market Interfaces (PMIs)** — buffer companies between the network and outside investors — open banking APIs, and jurisdiction-specific legal wrappers, intended to preserve network autonomy while remaining interoperable with external systems.

13.7 – Ethical and Regulatory Anchoring

Explains how MTUs embed the **ETHICAL values framework** into contracts, transaction flows, and governance processes. Regulatory compliance uses mission-locked legal design, privacy-preserving identity systems, and participatory auditability mechanisms.

13.8 – MTU as a Domain-Specific HAO Implementation

Concludes the section by presenting the MTU as a replicable financial implementation of the HAO model, not limited to the ICN. Unlike centralized banking and decentralized finance (DeFi) protocols, the MTU relies on relationship-based trust rather than credit scores or collateralized smart contracts, and is designed for interoperability with external financial systems while keeping reserve governance with members.

Key Contributions of the MTU Model

- **Financial Sovereignty:** Participants co-own, co-govern, and co-design the infrastructure they rely on.
- **Trust-Based Finance:** Capital access is determined by verified trust relationships rather than credit scores or interest-based lending.
- **Viability and Scale:** Designed to interface with existing financial systems while preserving cooperative governance.

- **Value Recirculation:** Surplus generated by network activity is reinvested in operating units and community infrastructure rather than distributed to outside capital.

The MTU model illustrates one way financial infrastructure can be organized around verified trust relationships and community-governed reserves rather than credit scores and centralized control.

Introduction

The **Member Trust Union (MTU)** is a domain-specific application of the **Humanized Autonomous Organization (HAO)** model addressing financial infrastructure, mutual credit, and community trust systems.

The **Integrated Cooperative Network (ICN)** defines the organizational topology and economic coordination of a distributed network of **United Micro Enterprises (UMEs)**. The MTU governs how capital flows, risk is shared, and value is stored, exchanged, and distributed among participants.

This section presents the MTU as a nested, polycentric, trust-anchored financial model that supports the ICN's operational and governance needs. Unlike traditional financial institutions, which rely on centralized authority, credit scores, and profit-maximization, the MTU is organized around relationship-based finance, intergenerational equity, and community-level governance of reserves. It maintains compatibility with external regulatory and financial systems, connecting grassroots value creation with formal capital systems.

Purpose of the MTU within the HAO/ICN Context

As the HAO Core governs the ICN's value systems and infrastructure coordination, the MTU governs the flow of financial trust, credit, and liquidity within and between UMEs. The MTU:

- **Allocates and recycles capital** at the edge, where value is created
 - Anchors the ICN's **trickle-out economic model** by directing shared capital outward, with separate diminishing-contribution and reinvestment pathways
 - Facilitates peer-to-peer microcredit, collective asset ownership, and localized financial resilience mechanisms
 - Provides a financial commons protected from high-interest external lending, capital leakage, and investment terms that would divert a majority of returns outside the network
 - Bridges the ICN with external banking systems and digital finance infrastructure through compliant and modular APIs, public market interfaces, and trust-scaling mechanisms
-

The MTU as a Humanized Financial Stack

Just as the ICN operationalizes a polycentric production network with reinvestment-based value flows, the MTU builds a financial stack around four concentric domains:

1. Individual-Level Trust Structures

Each member enters with progressive levels of verification and trust stewardship, contributing to peer-lending, local liquidity pools, and social collateral systems.

2. Local MTUs (LMTUs)

Hyper-local financial nodes embedded in communities, offering traditional and alternative financial services via hospitality-based physical spaces and relational finance models.

3. Regional Networks

Clusters of LMTUs that share standards, liquidity buffers, mutual aid protocols, and adaptive risk distribution frameworks.

4. Global MTU Infrastructure

Shared governance, compliance tooling, distributed ledger architecture, trust-verification systems, and connection points to external financial ecosystems (e.g. public credit unions, open banking APIs, and fintech rails).

Together, these layers are designed to provide financial access to members typically excluded from conventional banking, combining cooperative finance, mutual credit, and distributed-ledger tools within a single compliance framework.

MTU's Role in Reinforcing ICN Stability

As the financial anchor and trust substrate of the ICN, the MTU directly supports:

- **Liquidity access and reinvestment** for new and maturing UMEs
- **Compensation systems** aligned with value creation and equity distribution
- **Resilience buffers** via pooled reserves and emergency lending

- **Shared financial protocols** that standardize interactions between SEPs, UMEs, and external PMIs
- **Socioemotional stability**, intended to reduce the risk of member attrition during periods of economic scarcity

Where the ICN supports innovation and new-venture formation, the MTU is intended to provide financial coherence, protection, and rebalancing across the lifecycle of ventures and contributors.

Section Objectives and Structure

This section describes how the MTU operates as a HAO instance in its own right, while maintaining tight coupling to ICN operations. The following subsections (13.1–13.8) are outlined in the Executive Summary above.

Conclusion of Section Introduction

The MTU functions as a co-equal infrastructure layer to the ICN’s operational logic, not a secondary financial layer. In this model, trust replaces credit, reciprocity replaces interest, and network-level coordination replaces central banking. The MTU is intended to show how financial systems can be reoriented toward member well-being, long-term stewardship, and decentralized resilience.

§13.1 Multi-Level Trust Architecture

The **Member Trust Union (MTU)** — the network’s credit-union-like financial institution — operates as a layered, polycentric trust infrastructure that governs the flow of capital, risk, and responsibility without relying on traditional credit systems. The MTU treats trust as a primary design primitive, embedded in interactions, capital movements, and governance decisions, rather than as a binary or reputational overlay.

This architecture is composed of four interlocking layers, each with distinct roles in building, scaling, and anchoring trust across the network. Each layer is designed to reduce systemic fragility, increase local autonomy, and provide boundaries for risk containment.

A. Individual-Level Trust Structures

1. Progressive Trust Onboarding

Each member begins their MTU relationship through a progressive verification model:

- **Phase 0 – Intent Registration:** Declaration of values, basic identity confirmation, and passive observership.
- **Phase 1 – Relational Vouching:** At least two existing members must co-sign a social trust attestation.
- **Phase 2 – Transactional Trust:** Member engages in low-risk financial actions (e.g., micro-lending, rotating savings) and builds a transaction graph.
- **Phase 3 – Stewardship Rights:** Eligible to co-manage trust pools, initiate peer loans, and participate in LMTU governance.

Ostrom (2010) supplies a collective-action premise concerning trust and reputation. Narayanan et al. (2016) describe decentralized cryptographic identity, shared control, and limitations of reputation systems. The relational-credit stages and their proposed alternative to FICO-based scoring are corpus designs, not findings attributed to either source.

2. Trust Graph & Social Collateralization

Rather than relying on credit-scoring algorithms, MTUs use a multi-dimensional trust graph:

- Nodes = members; Edges = verified relationships (vouching, repayment history, co-investment, co-membership)
- Weighted edges encode trust depth, recency, and context
- Graph analytics provide contextual trust inference rather than deterministic scores

Loans, investments, and credit allocations are collateralized through shared trust rather than assets. In the event of default, trust erosion is distributed along relational paths, which is intended to deter overextension and encourage mutual due diligence.

B. Local MTUs (LMTUs)

1. Community-Embedded Finance Nodes

LMTUs serve as community-scale financial commons, offering:

- Pooled credit and liquidity buffers
- Mutual aid reserves

- Support for local UMEs (small, self-managing venture teams, $\leq \sim 15$ people) and SEP (a joint venture between teams) co-financing

Their governance is polycentric, often using consent-based decision-making or tiered circles based on participation level.

2. Hospitality-Based Physical Architecture

The corpus proposes community-bank and hospitality-inspired physical spaces for LMTUs. The following elements are design proposals:

- Walk-in financial commons (credit circles, group deliberation spaces)
- Shared-use community spaces (co-learning, resource exchange, mutual support)
- Conflict resolution rooms, advisory support from elder stewards

This design treats physical presence as a factor in trust-building, offsetting the limits of purely digital finance.

C. Regional Networks

1. Liquidity Federation & Mutualization

Clusters of LMTUs form regional federations to:

- Share liquidity buffers
- Pool and reallocate risk during economic shocks
- Coordinate distributed reserve targets based on predictive analytics (e.g., local employment trends, UME activity)

Each region operates under shared governance standards and a versioned **Regional Financial Protocol Agreement (RFPA)**.

2. Adaptive Trust Flow Protocols

Trust flows between LMTUs are governed by:

- **Threshold Trust Models:** Regions define their minimum relational requirements before extending pooled credit access
 - **Cross-Network Reputation Channels:** Allow high-trust individuals to migrate between LMTUs with a portable trust context
 - **Event-Triggered Failover Protocols:** Dynamic containment of trust erosion in one LMTU to prevent contagion (e.g., sudden UME collapse)
-

D. Global MTU Infrastructure

1. Meta-Ledger and Governance Stack

At the global level, the MTU maintains:

- A **meta-ledger** for anchoring local and regional trust states using cryptographic proofs rather than surveillance
- **Decentralized governance** for updating global rules (trust weighting algorithms, compliance schema, audit mechanisms)
- Federation of auditors and compliance anchors selected randomly, with consent override mechanisms, to balance transparency and autonomy

2. External Interface Layer

The global MTU layer maintains:

- **API gateways to external banking rails**, compliant with open banking standards (e.g., PSD2, FDX, ISO 20022)
 - **Value translation layers** that convert relational finance data into standardized reporting formats for funders, regulators, or aligned partners
 - **Permissioned transparency channels** for public market interfaces (e.g., Contribulo-like entities) to validate network creditworthiness without compromising local autonomy
-

Trust as a Dynamic, Multi-Sovereign Protocol

The MTU's trust architecture is nested and fractal rather than hierarchical:

- Each node (individual or LMTU) has sovereign trust-building capacity
- Trust flows dynamically through consent-based relational infrastructure
- Risk is contained through bounded autonomy and trust-dampening protocols
- Governance is adaptive, versioned, and locally overridable

This architecture is designed to support a plurality of trust cultures, adapt to local economic conditions, and operate at scale, in contrast to both centralized banking and purely algorithmic credit systems.

Conclusion of 13.1

By embedding trust directly into the structural design, across identity, liquidity, governance, and inter-node relationships, the MTU is intended to integrate closely with the social fabric of communities without sacrificing accountability or scalability. The multi-level trust architecture functions as both a financial substrate and a relational scaffold for the ICN (the reference cooperative business network), aligning individual action with collective resilience.

References

- Narayanan, A., Bonneau, J., Felten, E., Miller, A., & Goldfeder, S. (2016). *Bitcoin and cryptocurrency technologies: A comprehensive introduction*. Princeton University Press.
 - Ostrom, E. (2010). *Beyond markets and states: Polycentric governance of complex economic systems*. *American Economic Review*, 100(3), 641–672. <https://doi.org/10.1257/aer.100.3.641>
 - Scott, J. C. (1998). *Seeing like a state: How certain schemes to improve the human condition have failed*. Yale University Press.
 - Scholz, T., & Schneider, N. (Eds.). (2016). *Ours to Hack and to Own: The Rise of Platform Cooperativism, a New Vision for the Future of Work and a Fairer Internet*. OR Books.
 - McMillan, J., & Woodruff, C. (1999). *Interfirm relationships and informal credit in Vietnam*. *The Quarterly Journal of Economics*, 114(4), 1285–1320. <https://doi.org/10.1162/003355399556278>
 - Pentland, A., Lipton, A., & Hardjono, T. (2021). *Building the New Economy: Data as Capital*. MIT Press.
 - Ito, J., & Howe, J. (2016). *Whiplash: How to survive our faster future*. Grand Central Publishing.
-

§13.2 Financial Product Design in the Member Trust Union (MTU)

The MTU (the network’s credit-union-like financial institution) does not offer *financial products* in the conventional sense. Instead, it defines a set of financial protocols and social instruments embedded within a trust-centric infrastructure. These instruments are designed to facilitate circulation, cooperative ownership, mutual aid, and surplus reinvestment, while remaining interoperable with conventional financial systems.

This section outlines the architecture, principles, and mechanisms that make up the MTU’s financial design stack. Rather than services provided to members, these are primitives built into the ICN (the reference cooperative business network)’s operating structure. They are purpose-aligned, modular, and governed by MTU members across individual, local, and regional levels.

A. Design Principles

All MTU financial instruments adhere to six foundational principles:

1. **Purpose over Product**
Each instrument is designed around a *contextual need* (e.g., shared equipment, emergency liquidity, investment across UMEs (small, self-managing venture teams, $\leq \sim 15$ people)) rather than predefined “product classes.”
 2. **Trust-Centered Logic**
Instruments activate or scale in proportion to verified trust depth, not external collateral, formal credit scores, or capital exposure.
 3. **Reciprocity-Driven Returns**
Gains from instruments (interest-equivalents, surplus, or use-value) are directed into collective pools rather than distributed to individual holders.
 4. **Transparency by Default**
All instruments are governed by contract templates and collective audit tools accessible at the LMTU level.
 5. **Integration Across Scales**
Instruments function both within and across LMTUs, with clear protocols for portability, migration, and risk redistribution.
 6. **Surplus Reinvestment**
Instrument design directs value back into community infrastructure or UME development rather than into external capital markets.
-

B. Core Financial Instruments

1. Peer-Lending Pools (PLPs)

Overview:

PLPs are dynamically structured microcredit systems that allow member-to-member lending, trust-based underwriting, and liquidity provisioning without formal intermediaries.

Mechanisms:

- Members stake into rotating credit pools governed by relational trust depth.
- Credit access is granted through consent of the pool's trust circle (2–5 members), logged via a distributed trust ledger.
- Interest (if any) is expressed as reciprocal obligation rather than profit margin (e.g., lending labor, assets, or expertise back into the system).

Governance Features:

- Real-time dashboards for outstanding obligations and relational trust flows
 - Collective dispute resolution protocols for failed repayment
 - Adaptive pool ceilings based on LMTU liquidity and default rates
-

2. Collective Asset Finance (CAF)**Overview:**

CAF instruments allow multiple members or UMEs to co-own large assets (e.g., machinery, real estate, digital tools), either permanently or on use-based schedules.

Mechanisms:

- Asset ownership is fractionalized via tokenized legal wrappers or multi-party contracts
- Use is governed by time-share algorithms, cooperative calendars, or DAO-style voting
- Depreciation and maintenance costs are split based on usage, not capital contributed

Examples:

- Community-owned delivery vehicles used by multiple food or logistics UMEs
- Co-owned equipment libraries for prototyping, fabrication, or media production
- Shared commercial kitchen facilities managed via a usage-rights ledger

Economic Outcomes:

- Reduces redundant capital investment
 - Supports inter-UME collaboration
 - Allows long-tail contributors to access expensive assets through group trust
-

3. Mutual Aid Credit Instruments (MACIs)**Overview:**

MACIs function as embedded mutual aid protocols for emergency lending, life events, or short-term systemic shocks.

Mechanisms:

- Members opt into risk-pooling circles with predefined disbursement protocols (e.g., sickness, eviction, UME insolvency)
- Circle membership requires trust verification + contribution history
- Disbursements are activated by consent + algorithmic conditions (e.g., automated thresholds for cash reserves, distress indicators)

Trust Enforcement:

- Transparency of circle usage and surplus status
- Ability for contributors to adjust contribution tiers based on surplus risk

Institutional Precedents:

- Inspired by ROSCAs (Rotating Savings and Credit Associations), tontines, and kibbutz-like redistribution models
ROSCAs provide one rotating-contribution precedent (Ardener & Burman, 1995); savings groups provide a distinct accumulating-fund and share-out precedent (Gash & Odell, 2013). Tontines, kibbutz-like redistribution, and the MACI implementation are corpus synthesis rather than claims attributed to those sources.
-

4. Purpose-Aligned Investment Vehicles (PAIVs)

Overview:

PAIVs enable trust-anchored capital investment into UMEs or SEPs (joint ventures between teams), while limiting outside parties' capture of returns or control.

Mechanisms:

- Investments are structured as redeemable revenue-share agreements governed by MTU protocols, not venture-style equity
- Capital inflows may come from:
 - Reinvested MTU surplus
 - SEP-level capital raises
 - **Public Market Interfaces (PMIs)** — buffer companies between the network and outside investors — operating under MTU-aligned covenants
- Returns are redistributed based on:
 - Repayment triggers (e.g., % revenue thresholds)
 - Contributor classes (e.g., members, workers, adjacent LMTUs)
 - Diminishing Contribution logic (similar to the ICN)

Reinvestment Cycle:

- Investment → Local UME success → Partial return to PAIV → Rollover into next UME → Compounded collective gain

5. Liquid Commons & Community Treasury Protocols

Overview:

Each LMTU maintains a multi-tiered treasury for operational liquidity, long-term reserves, and targeted reinvestment.

Design Components:

- **Reserve Pool:** 3–6 months of operating capital, governed by collective thresholds
- **Commons Fund:** Allocated for community-benefit projects (e.g., free clinics, learning nodes, public infrastructure)
- **Risk Pool:** Redundancy mechanism triggered during macroeconomic shock or regional MTU failure

Mechanisms:

- Fund allocation governed via **Deliberative Budgeting Protocols (DBPs)**
- Treasury balance and allocation publicly auditable at LMTU and regional levels
- Minimum transparency standards required to interoperate with MTU federation

C. Interoperability with Conventional Financial Systems

While MTU instruments aim to shift economic logic, they are designed to interoperate with traditional financial systems through:

- **Open Banking APIs:** Integration with standard PSD2/FDX-compatible gateways for visibility, fiat exchange, and third-party auditing
- **Tokenized Contract Wrappers:** Legal templates that render MTU instruments legible to external actors (e.g., housing regulators, grantmakers, public funds)
- **Embedded Compliance Modules:** KYC/AML, identity verification, and fraud detection layered *behind* trust-based onboarding, triggered only when required for interoperability

This maintains a dual-flow system: MTU members can participate in traditional financial systems without adopting their underlying logic, and external entities can engage with MTU vehicles without altering MTU governance or internal values.

Conclusion of 13.2

The MTU's financial design departs from the standard logic of productization and credit abstraction, offering an architecture of financial primitives grounded in community needs, relationships, and cooperative governance. Through peer lending, collective asset use, mutual aid, and reinvestment-based financing, MTUs provide infrastructure intended to circulate trust as value and extend financial inclusion and resilience at the community level.

References (APA Style)

- Ardener, S., & Burman, S. (Eds.). (1995). *Money-Go-Rounds: The Importance of Rotating Savings and Credit Associations for Women*. Berg.
- Gash, M., & Odell, K. (2013). *The Evidence-Based Story of Savings Groups: A Synthesis of Seven Randomized Control Trials*. The SEEP Network.
<https://www.findevgateway.org/sites/default/files/publications/files/mfg-en-paper-the-evidence-based-story-of-savings-groups-a-synthesis-of-seven-randomized-control-trials-2013.pdf>

- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.
- Scholz, T., & Schneider, N. (Eds.). (2016). *Ours to Hack and to Own: The Rise of Platform Cooperativism, a New Vision for the Future of Work and a Fairer Internet*. OR Books.

§13.3 Governance and Risk Distribution in the Member Trust Union (MTU)

The MTU (the network's credit-union-like financial institution) combines financial instruments with polycentric governance logic and multi-scalar risk distribution mechanisms. Together these are designed to function as a decentralized financial infrastructure layer, capable of managing volatility, conflict, and uncertainty without collapsing into centralization or chaos.

In contrast to traditional financial institutions, which centralize risk at the top and offload failure to members (e.g., via overdraft fees, interest traps, or foreclosure), the MTU distributes risk horizontally and vertically across trust-based governance layers. It does this using consent-driven governance protocols, bounded liability pools, dynamic trust-weighting, and cascading containment systems that reflect both the social topology and the economic flow of the network.

A. Governance Architecture: Polycentric and Consent-Based

1. Fractal Governance Layers

The MTU operates as a nested, recursive governance network with the following strata:

- **Member-Level Micro Governance**
 - Personal trust stewardship
 - Peer lending consent
 - Initiation and oversight of mutual aid requests
- **Local MTU (LMTU) Circles**
 - Treasury allocation decisions (Deliberative Budgeting)
 - Resolution bodies for local financial disputes
 - Onboarding/vetting of financial products
- **Regional Federations**
 - Coordinated liquidity balancing
 - Emergency response coordination
 - Governance audits and protocol versioning
- **Global MTU Assembly**
 - Maintains shared financial operating protocols
 - Oversees interoperability with the ICN (the reference cooperative business network), public interfaces, and external regulators
 - Operates a rotating council of **Cross-Domain Stewards** (e.g., legal, technical, ethical domains)

Key Principle: Decision-making authority remains as close to the edge as possible while enabling network-wide coherence through explicit delegation.

2. Consent-Based and Trust-Weighted Voting

Rather than simple majority votes, the MTU applies:

- **Consent thresholds** (as in sociocracy): A proposal is adopted if no one has a reasoned, substantive objection.
- **Trust-weighted influence models**: Participation in high-risk or high-leverage decisions is calibrated by one's relational trust depth in the network (not tenure or wealth).
- **Deliberative Pre-Vote Protocols**: All financial decisions over a defined threshold must pass through structured dialogue phases: information → clarification → proposal → deliberation → decision.

B. Trust-Layered Risk Distribution

1. Risk as a Trust-Bound Field

Risk in the MTU is distributed based on the structure and strength of trust relationships rather than financial modeling alone. This builds on two key assumptions:

- Trust is contextual and multi-dimensional: risk in one domain (e.g., liquidity failure) is not identical to another (e.g., stewardship failure).
- Trust functions as a form of social capital intended to absorb uncertainty that financial capital alone may not cover.

2. Credit Containment Zones (CCZs)

To avoid cascading failure, MTUs define **Credit Containment Zones** at three levels:

- **Individual Zones:** Maximum credit exposure is constrained by individual trust boundaries.
- **Pool-Level Zones:** Mutual aid or lending pools define strict max-loss thresholds and shared guarantees.
- **Network Zones:** Regional MTUs enforce ceiling rules and auto-triggers to halt cross-region lending during volatility spikes.

These zones are dynamic, recalculated regularly based on:

- Trust erosion (missed obligations, peer feedback)
 - Economic signals (revenue drops, inflationary stress)
 - Changes to governance alignment or value adherence (via Value Alignment Monitoring (VAM) — ongoing checks that actions match stated principles)
-

C. Multi-Pool Liquidity and Shock Buffering

The MTU applies liquidity multi-pooling, meaning no single treasury or loan structure is system-critical. Liquidity is divided across:

- **Operational Pools** (daily access, transaction smoothing)
- **Stabilization Reserves** (emergency response, asset protection)
- **Mutual Aid Pools** (crisis lending, no expected return)
- **Reinvestment Pools** (strategic growth and capacity expansion)

Each pool:

- Operates with independent governance, tailored to risk appetite and member profile
 - Is monitored through predictive liquidity stress testing models (e.g., transaction graph modeling, volatility signals from SEPs (joint ventures between teams))
 - Has its own risk-tolerance profile, with clear documentation visible to members and auditors
-

D. Adaptive Trust and Credit Scoring Mechanisms

The MTU does not use static credit scores. Instead, it maintains a dynamic, contextual trust index for each member, pool, and UME (a small, self-managing venture team, $\leq \sim 15$ people).

1. Relational Trust Graph (RTG)

- Trust is recorded as interpersonal edges: e.g., repayments, co-investments, conflict resolution participation
- Weights include recency, depth, cross-domain exposure, and risk context
- Decay functions lower trust score over time in absence of activity

2. Trust-Weighted Credit Scaling (TWCS)

- Available credit scales according to the following factors, rather than income or asset value:
 - Number of active trust edges
 - Quality of relationships (measured through multi-party confirmations)
 - Participation in governance and mutual aid
 - When risk rises (e.g., regional liquidity shock), TWCS automatically reduces exposure ceilings and reroutes flows to members with stronger trust standing
-

E. Conflict Resolution and Risk Recovery Protocols

To limit the risk that conflicts erode network integrity:

- **Layered Resolution Paths:** Local → Regional → Global escalation (only if required)
- **Restorative Justice Model:**
 - Conflict circles convened
 - Trust reparation plans (service, education, etc.)
 - Time-bound probation instead of punitive bans
- **Failure Recovery Templates:**
 - Pre-agreed workflows for trust or financial collapse

- Allows restructuring or reallocation of obligations without total member removal
-

F. Regulatory Buffering and Legal Risk Channels

While financial decisions remain trust-driven, the MTU includes boundary infrastructure to satisfy external regulation without compromising its internal principles:

- **Audit-safe Wrappers:** Pools and agreements are wrapped in pre-approved, jurisdiction-specific legal formats (DAO LLC, coop corp, community benefit entity)
 - **Legal Risk Pools:** Collective insurance against regulatory penalties or violations
 - **Discretionary Transparency Protocols:** Regulatory observability is opt-in and audit-triggered, not always-on
-

G. Risk as Regenerative Capacity

The MTU's risk architecture treats risk as a signal that can be used to rebuild trust systems, in addition to mitigating failure:

- Failure in one domain often leads to protocol evolution in another (e.g., a lending collapse informs pool governance rules)
 - Trust erosion leads to community investment in education or conflict repair, rather than punishment
 - Members build risk literacy and collective responsibility as core financial competencies
-

Conclusion of 13.3

The MTU's governance and risk model departs from traditional assumptions of centralized underwriting, actuarial rigidity, and opaque risk transfer. Instead, it functions as an adaptive governance structure in which trust is measured, negotiated, repaired, and scaled as a dynamic field of accountability.

By rooting risk in relationships and responsibilities, in addition to returns and reserves, the MTU is designed to support community-scale financial stability, with systems intended to resist collapse and adapt after failure.

References (APA Style)

- Ostrom, E. (2005). *Understanding Institutional Diversity*. Princeton University Press.
 - Kostakis, V., & Bauwens, M. (2014). *Network Society and Future Scenarios for a Collaborative Economy*. Palgrave Macmillan.
 - Scholz, T., & Schneider, N. (Eds.). (2016). *Ours to Hack and to Own: The Rise of Platform Cooperativism, a New Vision for the Future of Work and a Fairer Internet*. OR Books.
 - Hardin, G. (1968). *The Tragedy of the Commons*. *Science*, 162(3859), 1243–1248. <https://doi.org/10.1126/science.162.3859.1243>
 - Lietaer, B., Arnsperger, C., Goerner, S., & Brunnhuber, S. (2012). *Money and Sustainability: The Missing Link*. Triarchy Press.
 - Poteete, A. R., Janssen, M. A., & Ostrom, E. (2010). *Working Together: Collective Action, the Commons, and Multiple Methods in Practice*. Princeton University Press.
-

§13.4 Technology and Physical-Digital Integration in the MTU

The Member Trust Union (MTU) — the network's credit-union-like financial institution — is not structured as a traditional bank or a digital-first fintech stack. Its technology functions as supporting infrastructure rather than the primary member-facing experience, coordinating human-led, trust-centered financial interactions. The MTU integrates physical space design, digital trust protocols, and data flows to support community-scale finance that is legible to members and interoperable with external systems.

This section describes the MTU's infrastructure stack. Rather than a layered software architecture, the stack is concentric: people are at the core, and digital tools support social processes rather than replacing them.

A. Foundational Design Principles

1. Human-in-the-Loop Coordination

Financial operations are designed to support human relationships rather than replace them. Fully autonomous decisions are not used in high-risk or interpersonal contexts.

2. Embedded Transparency

Transactions, pool states, and governance records are accessible to members directly, rather than through opaque APIs or dashboards.

3. Contextual Modularity

Technology components are composable, allowing each Local MTU (LMTU) to adopt modules appropriate to its community size, needs, and technical capacity.

4. Physical-Digital Reciprocity

Each digital capability has a corresponding physical interface in MTU hubs, intended to support accessibility and cultural fit.

B. MTU Technology Stack (Overview)

The MTU uses a set of technical systems that align with internal governance needs and external interoperability standards.

Layer	Function	Technology
Trust Graph Engine	Relationship tracking, vouching, trust-weighted interactions	Graph DB (e.g., Neo4j), custom trust metrics engine
Smart Trust Contracts	Programmatic enforcement of pooled agreements, trust-scaling, lending terms	DAO-like protocol with legal wrappers; WASM-based contract execution
Distributed Ledger Substrate	Anchor point for transaction proofs, trust state attestations	Private or federated DLT (e.g., Hyperledger Besu / Tendermint)
Consent OS	Governance workflow, participatory voting, deliberation logging	Custom modular framework; integrates Loomio-like patterns
Liquidity Pool Manager	Treasury state, stress-testing, capital routing	Simulation + accounting hybrid layer
Audit & Risk Analytics Engine	Value alignment monitoring, compliance visibility	Predictive models trained on trust graph + UME metrics
External Interface Bus	API gateway for integration with credit unions, regulators, or fintech partners	OAuth2 / OpenBanking APIs; optional FHIR-style schemas for financial records

C. Physical Infrastructure Design

Rather than building “branches” or call centers, MTUs use hospitality-centered spaces built around:

1. Relational Zones

- **Open Commons:** Where members meet, learn, plan lending circles, or explore new UME (small, self-managing venture team) ideas
- **Resource Rooms:** Semi-private spaces for tool sharing, co-creation, and transaction support
- **Stewardship Tables:** Analog governance stations—physical forms of deliberation and consent (e.g., color-coded tokens, trust-chain chains)

2. Consent Terminals

- Wall-mounted or mobile kiosks allow members to:
 - Review shared budget proposals
 - View mutual aid requests in queue
 - Swipe-to-consent on trust-linked decisions
- Paired with public projection surfaces showing real-time treasury health, governance items, and trust flow summaries

3. Sensory Trust Interfaces

To bridge abstract digital processes with embodied cues:

- Lighting gradients to indicate surplus/deficit status
- Haptic touch feedback on consent devices (e.g., slight pulse on decision activation)
- Audio cues for trust-confirmation events (e.g., musical tones for approvals, poetic snippets for rejections)

This use of physical and sensory cues, informed by service design and ritual theory, is intended to make community financial activity perceptible through more than visual display alone.

D. Trust Identity and Credentialing System

The MTU uses a **Progressive Trust Credentialing System** instead of standard logins or KYC:

Features:

- **Non-transferable digital credentials** based on verified relationships, not formal documents alone
- **Progressive disclosure:** only necessary data is surfaced per context
- **Credential tiers:** Used to gate access to higher-trust instruments (e.g., co-governance, pool formation)

Technically implemented using **Self-Sovereign Identity (SSI)** models (e.g., Verifiable Credentials, Decentralized Identifiers) with a local-first data posture (data remains with member devices unless explicitly shared).

E. Risk, Failure, and Redundancy Systems

1. Trust-Aware Redundancy

- Each digital system has a human failover path (e.g., paper-based consent during connectivity outages)
- Multi-node data propagation replicates ledgers, trust scores, and pool states across federated LMTU hubs

2. Deliberate Downtime Mode

- Periodic system pauses (e.g., one day per month) during which all MTU hubs operate in analog-only mode, intended to maintain interpersonal financial practice and preparedness for digital outages

3. Crisis Interoperability Mode

- Emergency financial flows (e.g., disaster relief) can be rerouted through inter-MTU bridges, with fallback liquidity caps and lockout periods intended to limit systemic risk
-

F. External Systems Compatibility

The MTU does not isolate itself from traditional financial systems; it encapsulates and regulates their influence:

- **Public Market Interfaces (PMIs)** — buffer companies between the network and outside investors — use a **Trust Compatibility Layer** to convert pooled MTU trust scores into formats legible to investors, banks, or regulators, without disclosing underlying member data
 - Integration with local credit unions, public benefit corporations, or municipal bonds happens through filtered APIs that translate intent as well as financial statements
 - Cross-border remittance is supported through stable-value bridge tokens (either CBDC-wrapped or community-issued scrip), intended to keep liquidity movement consistent with local trust relationships
-

G. Metrics and Feedback Systems

To maintain integrity and observability:

- Live trust flow maps, available to members and stewards
- Consent fatigue monitors flag when governance load may reduce engagement
- Sociometric dashboards aggregate feedback on space usage, deliberation quality, and emotional safety

All systems report to **Value Alignment Monitoring (VAM)** — ongoing checks that actions match stated principles — modules, which trigger protocol adaptation if misalignment is detected (e.g., surplus concentration, pool imbalances, governance bottlenecks).

Conclusion of 13.4

The MTU's technology operates in the background of member experience rather than at its center. Its infrastructure emphasizes physical presence and shared practice more than purely digital platforms do, with digital systems supporting trust-based interactions rather than replacing them.

The MTU's systems combine human-scale physical design, composable trust protocols, and direct transparency mechanisms, functioning as both a physical space and a technical platform.

References (APA Style)

- Greenfield, A. (2017). *Radical Technologies: The Design of Everyday Life*. Verso.
- Norman, D. A. (2013). *The Design of Everyday Things: Revised and Expanded Edition*. Basic Books.
- Ostrom, E. (2010). *Beyond markets and states: Polycentric governance of complex economic systems*. *American Economic Review*, 100(3), 641–672.
- Service Design Network. (2019). *Touchpoint*, 11(1), *Service Design for Innovation and Start-ups*. <https://www.service-design-network.org/touchpoint/vol-11-no1-service-design-for-innovation-and-start-ups>
- Pentland, A., Lipton, A., & Hardjono, T. (2021). *Building the New Economy: Data as Capital*. MIT Press.
- Bauwens, M., Kostakis, V., & Pazaitis, A. (2019). *Peer to Peer: The Commons Manifesto*. University of Westminster Press. <https://doi.org/10.16997/book33>

§13.5 Member Experience and Compensation Systems in the MTU

The Member Trust Union (MTU) — the network’s credit-union-like financial institution — is not a bank, credit union, or fintech company. It operates as a trust-centered financial commons, where membership is a layered, evolving relationship rather than a binary account status. This section outlines how members engage with the MTU across time, trust depth, contribution modes, and economic participation.

Rather than treating individuals as passive consumers of financial services, the MTU treats members as participants in value creation, capital allocation, and network resilience. Compensation combines baseline economic stability, cooperative performance participation, and long-term equity structures, intended to align personal financial stability with collective growth.

A. The Member Lifecycle in the MTU

Each member’s progression is not a fixed onboarding sequence. It follows trust-calibrated growth across six stages:

Stage	Name	Description
1	Initiate	Learns the MTU’s purpose, values, and basic rights/responsibilities. Gains initial access via peer sponsor or local orientation.
2	Participant	Actively uses MTU systems: peer lending, mutual aid pools, co-working spaces. Begins to contribute labor or governance input.
3	Contributor	Entrusted with responsibilities in local decision-making, capital circulation, or project leadership. Begins accruing equity.
4	Steward	Maintains high trust score. Participates in mentoring, complex deliberation, and pool management. Eligible for long-term compensation mechanisms.
5	Integrator	Bridges multiple MTUs or cross-network SEPs (joint ventures between teams). Guides trust calibration and helps resolve inter-node conflicts.
6	Elder	Non-active member with legacy equity, advisory rights, and symbolic governance presence. Holds cultural memory and ensures value continuity.

Progression is non-linear, with feedback loops, pauses, and re-engagement possible. Trust is re-earned, not permanently conferred.

B. Trust-Based Access, Not Tiered Privilege

Instead of traditional “premium” membership models, MTU instruments are accessed based on verified trust pathways:

- **Trust Edge Depth:** Number and quality of trust relationships (recorded via graph)
- **Reciprocal Contributions:** Prior work, resource sharing, governance participation
- **Situational Need:** Mutual aid and emergency access weighted toward members experiencing structural precarity

This is intended to create a permissioning model grounded in social relationships rather than economic stratification.

C. Compensation Architecture Overview

MTU member compensation is structured around four components:

Component	Purpose	Delivery Mode
1. Baseline Livelihood Compensation	Ensures minimum financial stability for active contributors	Weekly cash-equivalent distribution via pooled treasury
2. Performance-Linked Distributions	Rewards contribution to UME (small, self-managing venture team) or MTU-wide success	Monthly, based on pre-agreed local and network metrics
3. Network Profit Sharing	Redistributes surplus from MTU operations to members	Quarterly or annually, pro-rata via equity share class
4. Long-Term Equity Accrual	Builds intergenerational wealth and system co-ownership	Continuous, via tokenized or ledger-anchored equity in local node

These components are configured by each Local MTU (LMTU) and calibrated to regional economic conditions, under a shared governance agreement for network-wide harmonization.

D. Equity as Time-Bound, Trust-Weighted Ownership

1. Equity Issuance

Members earn equity based on:

- **Hours contributed** (normalized by task class)
- **Outcome-linked project participation**
- **Capital risk taken**
- **Relational labor** (mentorship, care, dispute resolution)

Equity is non-voting, non-transferrable outside the MTU, and subject to trust-based vesting curves.

2. Dynamic Vesting and Recovery

- Equity vests over time, with decay or pause during prolonged disengagement
 - If a member leaves under conflict or malfeasance, equity is subject to a trust-weighted recovery process: part returns to the commons pool, part to a dispute reparation fund
-

E. Timebanking and Non-Monetary Compensation

Many MTUs integrate complementary time- or service-based compensation systems, including:

- **Timebank Credits:** 1 hour of skilled labor = 1 credit, exchangeable for other services (e.g., childcare, transportation, tutoring)
- **Mutual Recognition Tokens:** Symbolic or fungible representations of appreciation, mentorship, or solidarity
- **Social Dividend Shares:** Earned for high-trust, low-visibility work (emotional labor, event hosting, accessibility work)

These flows are recorded and validated in member trust ledgers, influencing future compensation and governance access.

F. Accessibility and Inclusion Mechanisms

To reduce exclusion based on ability, time availability, or digital access:

- **Flexible Contribution Modes:** Participation pathways for care work, accessibility advocacy, storytelling, or conflict mediation
 - **Non-Monetary Merit Accrual:** Contribution points issued for organizing, emotional labor, or community rituals
 - **Compensation Choice:** Members can choose between direct payouts, equity accrual, or community fund contributions
 - **Language + Interface Equity:** All compensation tools must be legible in multiple formats: plain language, local dialect, visual maps
-

G. Intergenerational and Legacy Value Systems

The MTU does not treat compensation as zero-sum or purely contemporary. It includes:

- **Elder Compensation Tracks:** For legacy members providing cultural continuity, system memory, or care work
 - **Memorial Trusts:** Directed contributions in memory of deceased contributors, sustaining their work through designated equity pools
 - **Child/Dependent Equity Proxies:** Allocations to future generations via guardianship structures within the LMTU
-

H. Ethics of Compensation

All compensation systems are subject to:

- **Transparency Audits:** Equity and compensation logs are publicly viewable (with pseudonymization options)
 - **Deliberative Review:** Members can propose shifts in compensation logic via governance deliberation
 - **Consent-Bound Compensation:** Compensation cannot be structured to draw value from another member's labor without that member's consent and trust verification
-

Conclusion of 13.5

The MTU frames compensation as recognition of relational contribution and system co-stewardship, not solely as a reward for productivity. Members receive support for contributing to a shared commons, rather than being paid only for discrete work output.

The MTU replaces wage-based compensation with trust-aligned, equity-generating participation, changing how earning, membership, and long-term participation are structured within the financial system it operates.

References (APA Style)

- Scholz, T., & Schneider, N. (Eds.). (2016). *Ours to Hack and to Own: The Rise of Platform Cooperativism, a New Vision for the Future of Work and a Fairer Internet*. OR Books.
- Piketty, T. (2014). *Capital in the Twenty-First Century*. Harvard University Press.
- Gibson-Graham, J. K. (2006). *The End of Capitalism (As We Knew It): A Feminist Critique of Political Economy* (New ed.). University of Minnesota Press.
- Standing, G. (2011). *The Precariat: The New Dangerous Class*. Bloomsbury Academic.

- Weeks, K. (2011). *The Problem with Work: Feminism, Marxism, Antiwork Politics, and Postwork Imaginaries*. Duke University Press.
- Dardot, P., & Laval, C. (2019). *Common: On Revolution in the 21st Century*. Bloomsbury Academic.

§13.6 Integration with ICN and External Markets

The Member Trust Union (MTU) — the network’s credit-union-like financial institution — is not a standalone financial entity. It functions as a financial substrate of the Integrated Cooperative Network (ICN) — the reference cooperative business network. UMEs (United Micro Enterprises) — small, self-managing venture teams of up to about 15 people — create value, and the HAO (Humanized Autonomous Organization) — the network’s coordinating framework — coordinates governance and systems integration. The MTU is responsible for storing, distributing, exchanging, and protecting value according to trust-based principles.

This section describes how the MTU integrates with the ICN’s operational and governance systems, and how it selectively interfaces with external financial and regulatory infrastructure. These integrations are intended to let the MTU remain resilient while permeable to outside systems, and cooperative while legible to capital markets.

A. Internal Integration with ICN Systems

The MTU is not an add-on service to the ICN; it functions as a core infrastructural layer coordinated with other components of the HAO-based network.

1. UME Alignment and Resource Flow

- UMEs register with their Local MTU (LMTU) for financial participation (e.g., investment, compensation, pooled liquidity)
- MTUs assess UME trustworthiness via:
 - Enterprise Integration Assessments (EIA)
 - Value Alignment Monitoring (VAM) — ongoing checks that actions match stated principles
 - Reputation graphs spanning collaborations formed through SEPs (joint ventures between teams)
- Once accepted, a UME gains access to:
 - Operational liquidity pools (with risk-limited ceilings)
 - Peer-to-peer credit
 - Collective asset financing
 - Revenue routing systems with diminishing contribution schedules

2. SEP Financial Routing and Multiparty Flow Logic

- When UMEs form Strategic Enterprise Partnerships (SEPs), a composite trust entity is created
- MTUs define multi-party flow contracts, governing:
 - Shared liabilities
 - Contribution-weighted returns
 - Pooled reinvestment rates
- Flow contracts are anchored in the MTU ledger and can evolve based on SEP performance metrics (e.g., internal efficiency, external revenue, alignment)

3. Governance Layer Coupling

- MTU stewards are often the same individuals involved in HAO-level strategy circles
 - MTU financial protocols respond to changes in the DEA (Dynamic Enterprise Agreement) — a versioned operating agreement replacing fixed bylaws — and provide signal feedback (e.g., liquidity stress, surplus zones)
 - Consent OS systems used in MTUs and UMEs are federated and interoperable, enabling shared governance pathways
-

B. Reinvestment and Redistribution Logic

MTUs implement the trickle-out economic model by channeling capital outward:

- Downward for operational liquidity (to UMEs and members)
- Laterally for SEP development

Diminishing contributions return upward to shared HAO infrastructure, then are allocated outward again as reinvestment or surplus allocation.

This system relies on:

- **Scheduled revenue recirculation:** Monthly, quarterly, annual redistributions

- **Performance-modulated flows:** Adjusted by each UME's contribution to shared goals
 - **Equity-balancing protocols:** Used to limit overaccumulation in high-performing nodes while maintaining autonomy
-

C. External Integration with Financial Systems

While grounded in cooperative finance, the MTU engages selectively with external infrastructure. This occurs via Public Market Interfaces (PMIs) — buffer companies between the network and outside investors — and compliance-oriented components.

1. Public Market Interfaces (PMIs)

PMIs function as intermediary zones between the ICN and outside capital markets. They are:

- Purpose-built SEPs (e.g., Contribulo-like entities)
- Legally structured (e.g., DAO LLCs, Public Benefit Corporations)
- Governed by multi-party trust contracts rooted in ICN and MTU principles

PMIs provide:

- Access to outside investment without compromising internal sovereignty
- Regulatory measures that buffer UMEs from profit-maximizing investor behavior
- Convertible instruments (e.g., redeemable equity, mission-locked shares)

2. Open Banking and Payment Rails

MTUs integrate with external banking services via selectively permeable APIs, which include:

- OpenBanking-compliant access points for read/write financial data
- Programmable payment channels (e.g., via stablecoin bridges or real-time gross settlement rails)
- Sovereign identity protocols (e.g., DID + verifiable credentials) for interfacing with KYC/AML-compliant systems without full data exposure

3. Regulatory Compatibility

MTUs create legal wrappers for their pools, contracts, and instruments, such as:

- Cooperative corporations
- Federated mutual associations
- Jurisdictional DAO LLCs
- Local community development finance institutions (CDFIs)

These wrappers allow MTUs to:

- File taxes
 - Issue certified instruments (e.g., revenue-backed notes)
 - Offer FDIC-insured equivalents (via proxy partnerships with credit unions or state banks)
-

D. Risk Isolation and Containment for External Access

To limit contagion or value leakage:

- All external access is sandboxed behind capital firewalls
- External inflows are converted into commons-wrapped instruments (e.g., mutual credit tokens, redeemable bonds)
- External exits require multi-party consent from internal stewards and affected members

This selective permeability is intended to keep external interaction aligned with commons priorities rather than outside interests focused primarily on capital return.

E. Strategic Value of Interoperability

By remaining interoperable with external systems without being dependent on them, MTUs provide the ICN with:

- **Resilience:** Ability to continue operating during disruption to global financial systems

- **Legibility:** Optional reporting for compliance or public accountability
- **Funding Flexibility:** Ability to raise mission-aligned capital
- **Labor Mobility:** Portable trust scores across ICN and partner ecosystems

F. Diagram: Dual Integration Map (suggested)

A concentric diagram showing:

- Core MTU ↔ HAO and UMEs (internal)
- MTU edge ↔ PMIs, external banks, investors (external)
- Arrows with gating logic, trust weighting, and compliance filters

Conclusion of 13.6

The MTU’s integration design is intended to give the ICN financial infrastructure that combines the autonomy of a self-sufficient system with the interoperability of a platform-native institution.

Rather than isolating from capital markets or fully integrating with them, the MTU applies a selective-interface model: flow is permitted where values align and restricted where they diverge. This gives the ICN interfaces for adaptation and risk containment beyond those of a standalone cooperative federation.

 References (APA Style)

- Scholz, T., & Schneider, N. (Eds.). (2016). *Ours to Hack and to Own: The Rise of Platform Cooperativism, a New Vision for the Future of Work and a Fairer Internet*. OR Books.
- Lietaer, B., Arnsperger, C., Goerner, S., & Brunnhuber, S. (2012). *Money and Sustainability: The Missing Link*. Triarchy Press.
- Bauwens, M., Kostakis, V., & Pazaitis, A. (2019). *Peer to Peer: The Commons Manifesto*. University of Westminster Press. <https://doi.org/10.16997/book33>
- Ostrom, E. (2005). *Understanding Institutional Diversity*. Princeton University Press.
- Pentland, A., Lipton, A., & Hardjono, T. (2021). *Building the New Economy: Data as Capital*. MIT Press.

§13.7 Ethical and Regulatory Anchoring

The MTU — the network’s credit-union-like financial institution — is not ethically unbounded or legally unanchored, despite being a new financial infrastructure design. This section outlines how the MTU operationalizes ethical principles and maintains regulatory compatibility without compromising its human-centered, trust-based architecture.

The MTU embeds ethics into protocol design, legal structure, and cultural practice, in addition to complying with external regulations. This section details how MTUs meet fiduciary, compliance, and commons stewardship responsibilities while maintaining sovereignty.

A. Core Ethical Design Principles

The MTU adopts the ETHICAL framework of the HAO — the network’s coordinating framework. Under this framework, financial practices are intended to be mission-aligned, values-governed, and legible in social and emotional terms as well as financial ones.

ETHICAL Principle	MTU Operationalization
Empathy	Trust calibration models emphasize lived context over credit abstraction
Transparency	Open ledger views for all members, participatory budgeting, auditability by non-experts
Harmony & Holistic Health	Anti-scarcity design: liquidity flows structured for collective wellbeing, not profit-maximization
Integrity	Role-based accountability structures, traceable decisions, peer-verifiable governance
Collaboration	Consent-based agreements, pooled ownership, and deliberative governance
Accountability	Dynamic audit logs, social reputation systems, restorative justice practices
Learning & Longevity	Feedback loops, financial education integration, and elder stewardship roles

The MTU does not separate ethical policy from technical operations: values are enforced in code, contracts, and daily practice.

B. Embedded Ethical Safeguards

1. **Value Alignment Monitoring (VAM)** — ongoing checks that actions match stated principles
All financial transactions, flows, and governance decisions are subject to continuous VAM tracking:

- Cross-checked against community-voted principles
- Flagged for divergence (e.g., disproportionate resource concentration, exclusionary patterns)
- Triggers reflective governance dialogues rather than automatic punitive action

2. Consent Architecture

- No critical decision (e.g., external capital interface, pool deployment, SEP (joint venture between teams) restructuring) is taken without consent quorum
- Consent is informed, deliberative, and adaptive, drawing on sociocratic and deep democracy traditions

3. Ethical Gatekeeping for Innovation

- All new tools, interfaces, and market-facing instruments are subjected to an **Ethics Compatibility Review**
- Assessed against: member impact, systemic risk, regulatory implications, long-term cultural effect

C. Legal and Regulatory Anchoring

MTUs are structured to operate across jurisdictions and adapt to differing regulatory regimes, using a federated legal wrapper model:

Layer	Entity Type	Function
Local MTU	Co-op, Mutual Society, or CDFI-equivalent	Community governance, local compliance, fiduciary accountability
Regional Network	Multi-Stakeholder Cooperative or Umbrella Org	Shared liquidity and cross-node compliance handling
Global Coordination Layer	Public Benefit Corporation / DAO LLC	External interface, global licensing, audit and capital risk controls

Legal Features:

- **Mission Lock:** All MTUs include charter-based clauses that restrict external control, asset liquidation, or deviation from ETHICAL values.
- **Commons Clauses:** Certain IP, data, and processes are designated as non-alienable public infrastructure.
- **Distributed Fiduciary Responsibility:** Accountability is shared across governance roles, rather than centralized in directors or boards.

D. Regulatory Interface Models

Rather than evade compliance, MTUs translate compliance requirements into values-compatible protocols.

Examples:

- **AML/KYC Equivalents**
 - Uses verifiable credential systems (SSI/DID) to verify identity without disclosing sensitive information.
 - Trust path depth and endorsements used as dynamic risk scoring rather than static account limits.
- **Tax Compliance**
 - MTUs use flow accounting and smart trust contracts to automatically:
 - Allocate pooled taxes
 - Generate jurisdiction-specific reports
 - Support members in understanding and filing obligations
- **Consumer Protection**
 - Members can trigger **Collective Protective Protocols** to pause access, freeze pools, or initiate review panels if harm or misalignment is detected.
- **Financial Instrument Classification**
 - MTU-issued tokens are functionally designed to avoid being securities.
 - If tokenized equity is used, it is mission-locked and structured under regulatory exemptions or via certified cooperatives.

E. Sovereignty Safeguards

Despite external interfaces, the MTU maintains financial sovereignty and community autonomy through:

1. Selective Permeability Controls

- External capital may enter via PMIs (buffer companies between the network and outside investors), but is isolated from core pools
- External exit requires consent-based approval from affected stakeholders

2. Contributor-Restricted Returns

- No dividend, interest, or return can be issued to an entity that is not:
 - A verified contributor
 - Operating under commons-compatible rules

3. Commons Reclamation Protocols

- If an MTU experiences mission drift or capture:
 - Equity reverts to trust pools
 - Governance is rebalanced toward elders and regional stewards
 - Triggered via multi-level VAM and governance flag systems

F. Conflict Resolution and Restorative Mechanisms

All MTUs maintain dispute-resolution systems that combine:

- **Rapid Resolution Circles:** Peer panels with restorative orientation
- **Consent Withdrawal:** Any member may revoke consent to interactions (with consequences on both sides)
- **Reparative Equity Pools:** Used for harm remediation, funded from surplus rather than penalties charged to individual members

These systems are intended to reduce adversarial escalation and support long-term community resilience.

Conclusion of 13.7

The MTU is designed to be governed from within rather than regulated only from above. Regulatory compliance, member sovereignty, and ethical alignment are intended to operate together across its flows, protocols, and policies.

In the MTU, legal structure and value commitments are treated as complementary rather than opposed, intended to anchor the ICN (the reference cooperative business network)’s legitimacy in both regulatory compliance and commons accountability.

References (APA Style)

- Lessig, L. (2006). *Code: Version 2.0*. Basic Books.
- Ostrom, E. (2005). *Understanding Institutional Diversity*. Princeton University Press.
- Hardin, G. (1968). *The Tragedy of the Commons*. *Science*, 162(3859), 1243–1248. <https://doi.org/10.1126/science.162.3859.1243>
- Dardot, P, & Laval, C. (2019). *Common: On Revolution in the 21st Century*. Bloomsbury Academic.
- De Filippi, P, & Wright, A. (2018). *Blockchain and the Law: The Rule of Code*. Harvard University Press.

§13.8 MTU as a Domain-Specific HAO Implementation

The Member Trust Union (MTU) — the network’s credit-union-like financial institution — is a domain-specific implementation of the Humanized Autonomous Organization (HAO) model — the network’s coordinating framework — operating within the Integrated Cooperative Network (ICN) — the reference cooperative business network. It translates cooperative values into operational, scalable financial infrastructure.

The broader HAO framework provides a model for designing distributed, ethical, and resilient organizations. The MTU illustrates how these principles apply specifically to finance, capital flow, risk distribution, and trust-based coordination. This final subsection synthesizes the design patterns described throughout Section 13 and presents the MTU as a component of the ICN relevant to future cooperative economies.

A. HAO Characteristics Fully Expressed in the MTU

The MTU implements HAO architecture across each layer of its operation:

HAO Principle	MTU Implementation
Human Primacy	Trust replaces credit scores; human relationships anchor financial logic
Distributed Autonomy	Nested MTUs operate semi-independently, federated through interoperable agreements
Polycentric Governance	Consent-based decision-making at local, regional, and global levels
Regenerative Economics	Capital circulates rather than accumulates; wealth-building is participatory
Intentional Technological Integration	Seamless coupling of physical presence with digital finance infrastructure
Commons Stewardship	Surplus is reinvested in cooperative infrastructure; assets are collectively governed
Resilience and Adaptation	Multi-tiered liquidity pools, trust rebalancing, and mission-locks limit external claims on MTU assets

This alignment results from deliberate architectural design. The MTU functions as a mechanism for value exchange structured around interdependence, inclusion, and long-term continuity.

B. Financial Viability

1. Capital Efficiency and Flow Optimization

- The MTU allocates resources at the edge, in UMEs (small, self-managing venture teams), where value is created, rather than in central treasuries or fee-charging intermediaries.
- Redundant liquidity, high-trust peer lending, and cooperative reinvestment are intended to produce shorter capital cycles, lower volatility, and higher resource utilization.

2. Risk Distribution and Resilience

- Distributed mutual credit, diversified SEPs (joint ventures between teams), and dynamic risk weighting systems are intended to reduce the chance that a single failure destabilizes the whole system.
- MTUs are designed to contain financial shocks locally while drawing on shared regional reserves during crises, an approach informed by antifragile design in ecology and network theory.

3. Revenue Generation Capacity

- MTUs support monetization of local economies through:
 - Peer-to-peer lending interest redistribution
 - Community asset pools generating long-term rental or usage yield
 - Equity conversion interfaces with compliant public markets (PMIs, buffer companies between the network and outside investors)
 - These revenue models route returns to contributors and members rather than outside intermediaries, and are tied to member compensation systems.
-

C. Social Viability and Inclusivity

- MTUs are intended to lower barriers to financial participation by offering:
 - Alternative pathways to liquidity (social collateral, trust verification)
 - Consent-based economic engagement models
 - Timebanks and non-monetary exchanges for care work, mentorship, and other unpaid labor
- This is intended to expand financial inclusion, creating space for contributions often marginalized in traditional finance (care work, art, relational labor).

MTUs are also designed to be culturally legible: members can understand the logic behind the system, see how their value is measured, and participate in shaping financial tools and outcomes, narrowing the separation between users and the system that serves them.

D. Technological and Legal Scalability

- The MTU is built to be jurisdictionally modular and technologically extensible:
 - Integrates with legacy banking systems via API
 - Uses verifiable credentials and programmable compliance to meet KYC/AML without invasive surveillance
 - Legal wrappers (co-ops, DAO LLCs, benefit corps) allow participation in traditional systems without losing control
- As such, MTUs can scale across regions, federate into broader trust networks, and interface with traditional capital markets while retaining their own governance terms.

This allows the MTU to function as a bridge between emerging cooperative economies and legacy financial institutions.

E. Strategic Significance within the ICN

The MTU functions as an infrastructure layer relevant to the ICN's resilience, scalability, and ethical coherence:

- **Liquidity Routing Backbone:** All operational and reinvestment flows are channeled through MTU protocols, maintaining alignment with network intent.
- **Trust Continuity Engine:** MTUs preserve member, UME, and SEP reputations across time and geography, enabling long-term coordination.
- **Commons Integrity Safeguard:** MTUs enforce mission-locks, public benefit contracts, and cultural protocols intended to prevent drift, dissolution, or unauthorized capture.

- **Adaptability Lever:** Because MTUs are modular, they allow the ICN to enter new regions, partner with compliant institutions, and onboard diverse economies without departing from its principles.

The MTU is intended to let the ICN scale while preserving its founding principles, coordinating small local units within a larger federated structure.

F. Beyond the ICN: Model for a Cooperative Future

Although designed for the ICN, the MTU architecture could serve as a template for other cooperative or post-capitalist economies:

- Urban mutual aid groups could deploy localized MTUs for financial sovereignty.
- Federations of cooperatives could use MTUs to collectively fund and manage joint infrastructure.
- Global South solidarity networks could build MTUs to bypass international lending systems that impose high-interest, restrictive terms.

With appropriate legal scaffolding and digital tooling, MTUs could federate into interoperable global commons finance networks operating between state and market institutions.

Conclusion of 13.8

The Member Trust Union functions as the financial foundation of the ICN and as an instantiation of the Humanized Autonomous Organization model that could, in principle, be adapted elsewhere.

The MTU's multi-level architecture, consent-based governance, commons-preserving flows, and trust-based legitimacy structure are intended to offer a financial alternative to both centralized banking and speculative Web3 markets.

The MTU, as specified here, represents one operational model for pursuing those design goals within the ICN.

References (APA Style)

- Lietaer, B., Arnsperger, C., Goerner, S., & Brunnhuber, S. (2012). *Money and Sustainability: The Missing Link*. Triarchy Press.
 - Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.
 - Scholz, T., & Schneider, N. (Eds.). (2016). *Ours to Hack and to Own: The Rise of Platform Cooperativism, a New Vision for the Future of Work and a Fairer Internet*. OR Books.
 - Dardot, P., & Laval, C. (2019). *Common: On Revolution in the 21st Century*. Bloomsbury Academic.
 - Bauwens, M., Kostakis, V., & Pazaitis, A. (2019). *Peer to Peer: The Commons Manifesto*. University of Westminster Press. <https://doi.org/10.16997/book33>
-

§13.9 Physical Branch Model

The Member Trust Union (MTU) — the network's credit-union-like financial institution — has one existing description of its physical space: §13.4's Section C ("Physical Infrastructure Design"), roughly two dozen lines covering three zone types for a single generic hub — Relational Zones, Consent Terminals, and Sensory Trust Interfaces — plus a one-sentence note that the design draws on service design and ritual theory. It does not describe how a location is chosen, how it is staffed, how its space is scheduled across competing uses, or how physical and digital channels connect beneath the consent kiosks it names. This section supplies that layer: space design principles, a location-selection methodology, a staffing and competency framework, guidance for running one space across multiple uses, and the physical-digital touchpoints beneath the sensory interfaces §13.4C already describes. It is a companion to §13.4C, not a replacement for it; where the two sections describe the same feature from different angles, this section says so rather than restating the description.

Source and marker convention. The material below is drawn from a single staged extraction of six unreviewed source files, produced in one generative pass with no citations and never checked against a primary source. Roughly half of the underlying files present invented first-person quotations as field evidence; none of that testimony appears here, as quotation or as paraphrase implying a real speaker — passages that were purely quotational, with no separable design content, have been omitted rather than summarized. The source supplies almost no numeric detail. Across its full scope there is exactly one quantitative claim (§13.9D, below), and it appeared only inside a fabricated quotation; it is not reproduced here, because a figure whose sole provenance is invented testimony is fabricated data rather than an unchecked claim, and the marker (**unverified**) used elsewhere in this corpus would misdescribe it. Everywhere else, the source states a design category, a process, or a named list without attaching a measurement, and this section says so directly rather than supplying a plausible number. No section below carries a reference list: the source cites nothing, and supplying citations it does not have would manufacture a rigor it never had. Structural counts that describe how the source organizes its own material — four principles, six zones, five competency levels — are not measurements of anything in the world and are not separately marked; they describe the shape of an unreviewed design proposal, not a verified fact about it.

A. Space Design Principles

MTU space design is organized around four stated principles, framed throughout as a departure from traditional financial-institution design — imposing facades, formal interiors, and physical barriers between staff and members:

- **Hospitality over intimidation** — welcoming, low-anxiety environments; cultural relevance to the local community; accessibility across abilities and backgrounds.
- **Transparency and visibility** — layouts that make service processes visible to members; easy staff access; environmental design that communicates information directly rather than through signage alone.
- **Flexibility and adaptability** — reconfigurable areas; a design that scales across capacity levels; infrastructure built to evolve rather than be replaced.
- **Resource efficiency** — spaces sized to actual need rather than institutional display; sustainable-material selection; operational systems that minimize ongoing resource consumption.

The source gives no floor-area or square-footage figures, occupancy limits, or numeric zoning ratio for any of the above. Sizing is stated only as qualitative design intent, left to contextual design at each location.

Six spatial elements recur across MTU locations, adapted to context rather than applied as a fixed floor plan. The source gives no adjacency diagram and no required-adjacency rule connecting them; how they relate to each other in a given building is left to the design process described below.

Element	Function
Welcome Zone	Identifiable, accessible entry point; immediate orientation to what the space offers; seating and basic information on arrival.
Community Commons	The space's primary gathering area; flexible seating; displays of community and MTU information; supports spontaneous interaction.
Conversation Spaces	Semi-private areas with visual and acoustic separation for financial discussion and planning, plus fully enclosed rooms for confidential conversations; technology supporting remote participation.
Learning Environments	Workshop areas, self-directed digital-learning tools, a resource library, and hands-on demonstration space for financial concepts.
Operational Areas	Transaction service points, staff workspace, and back-of-house support space; the source calls for secure elements for sensitive activity without specifying what they are.
Digital Integration Points	Member-facing and staff-support technology, plus environmental interfaces that surface digital information within the physical space — see §13.9E, below.

MTU physical presence takes one of three forms, applied by context rather than a single template:

1. **Dedicated Facilities** — a purpose-built building, a renovated structure, leased dedicated space, or a mobile unit.
2. **Co-Located Services** — space shared with a community center, a social-enterprise hub, a public facility such as a library, or an educational institution.
3. **Hybrid Physical-Digital Presence** — minimal fixed space, using pop-up locations, staff who travel to members, partner-provided space, and a digital-first posture with targeted physical touchpoints.

Spaces are designed and evolved through four recurring practices: **participatory design** (design workshops, prototype testing, ongoing community feedback), **contextual adaptation** (cultural and climate fit, reuse of existing assets, explicitly framed as the alternative to a standardized branch template), **progressive implementation** (starting from available resources rather than an ideal end state, with defined expansion pathways), and **shared learning across the network** (documented design patterns, including approaches that did not work, shared across locations).

The design process itself runs in four phases:

Phase	Content
Community Engagement	Needs assessment, asset mapping, cultural exploration, partnership development with potential space partners.
Collaborative Design	Co-creation workshops with community and stakeholders; integration of design and financial-service expertise; prototype development and iterative refinement.
Implementation Support	Network-provided design guidance, expert consultation, and insight-sharing from comparable projects elsewhere in the network.
Continuous Evolution	Observation of actual usage, ongoing feedback collection, and adaptation cycles as community needs shift.

B. Location-Selection Methodology

Location selection is framed throughout the source as a departure from how traditional bank branches are sited — a process the source characterizes as driven primarily by market opportunity and profit potential. MTU siting instead applies four criteria and a four-phase process, and the source quantifies neither: it gives no numeric scoring rubric, no weighting across the four criteria, and no minimum-viability threshold — no minimum member count, no income figure, and no dollar threshold anywhere in the location-selection material. Where a scoring or weighting scheme would ordinarily go, the source names a category of analytical tool (“Viability Modeling,” below) without a formula.

The four selection criteria:

- **Community need prioritization** — financial-service gaps in the area; potential to improve local economic resilience; existing-member requests for physical service; explicit attention to demographic groups traditional institutions often overlook.
- **Existing-asset leveraging** — identifying underused space for repurposing; co-location with aligned organizations; integration with existing community gathering places.
- **Sustainability and viability** — long-term operational viability; a realistic read of the resources needed to establish and maintain the location; risk and resilience factors. The source states this criterion explicitly as not profit-driven, while still requiring the location to be sustainable.
- **Network considerations** — how a new location’s coverage complements the existing network of MTU locations; potential for resource-sharing with nearby locations; alignment with regional development strategy.

The selection process runs in four sequential phases:

Phase	Content
Initial Assessment	Geographic mapping of existing financial services and service gaps; demographic review; local economic conditions; relationship to the existing MTU network.

Phase	Content
Community Engagement	Structured conversations with residents and organizations; detailed needs assessment; collaborative asset mapping; discussions with potential local partners.
Site-Specific Evaluation	Physical characteristics and adaptability of candidate sites; accessibility analysis; cost projection for establishment and operation; regulatory review.
Collaborative Decision-Making	Detailed proposals, community feedback on specific proposals, structured comparison across options, and consensus-building toward a final selection.

Three location types, each with distinct siting criteria:

- **Urban Neighborhood Locations** — connection to a recognized neighborhood identity; public-transit accessibility; cultural alignment with the neighborhood; visibility and safety sufficient that members feel comfortable visiting.
- **Rural Service Hubs** — located at an existing community gathering point, within reasonable travel distance of the communities served (described only as centrally accessible; the source gives no mileage figure), with multiple-use potential and resource-efficient operation suited to limited infrastructure.
- **Special-Purpose Locations** — hubs for specific cultural or linguistic communities, financial-education centers, nodes supporting local economic-development initiatives, and sites used to test new physical-presence formats.

Three named categories of analytical tool, none with a stated formula:

- **Geographic Information Systems (GIS)** — service-gap mapping, demographic overlays, transportation and access analysis, and mapping of existing community assets.
- **Community Needs Assessment Toolkit** — quantitative survey instruments, qualitative methods, and participatory assessment techniques.
- **Viability Modeling** — cost-projection models, service-volume scenario planning, resource-requirement analysis, and identification of sustainability thresholds. The source names this last item only as a concept; no numeric threshold accompanies it anywhere in the material.

Short of a full permanent site, the source describes transitional presence (scheduled mobile-service visits, temporary pop-ups, phased establishment), technology-enhanced reach (digital services treated as complementary to physical presence, not a substitute), and community ownership models (mechanisms for local members to invest in their local space, with democratic decision processes for location decisions specifically).

Three implementation challenges and their named responses:

Challenge	Named response
Resource constraints	Phased development, resource pooling across stakeholders, in-kind contributions, efficiency innovations.
Competing location preferences among stakeholders	Transparent decision criteria, multi-benefit site analysis, complementary solutions, sequenced development over time.
Regulatory requirements	Regulatory mapping, compliance strategies, relationship-building with regulators, exploration of alternative legal structures where regulation blocks a specific approach.

The source states that the selection methodology itself is revised over time, through case-study documentation, periodic process review, and adjustment based on how prior location decisions performed. It gives no specific review cadence.

C. Staffing Model and Competency Framework

Staffing rests on four stated principles: **people over positions** (valuing a person’s full range of capability rather than a fixed job description, strengths-based rather than deficit-based), **community-connection priority** (local knowledge, cultural competence, and existing community relationships treated as staffing criteria in their own right), **balanced expertise** (four knowledge types treated as co-equal: formal professional and technical skill, lived experience, community wisdom, and diverse perspective), and **collaborative focus over hierarchy** (team complementarity, distributed leadership and accountability, mutual peer support).

Four staffing models, not mutually exclusive and applied by context:

Model	Composition
Core Team Model	A small core team in four defined roles — Local Guide (community-connected relationship building), Financial Navigator (financial expertise), Operations Coordinator (keeps the space and its systems running), and Learning Facilitator (supports financial-capability development) — often supplemented by specialized contributors.
Community Hybrid Model	A small number of consistent staff anchors, plus member contributors in defined supporting roles, specialized community volunteers, and personnel from partner organizations.
Network Support Model	Used for smaller locations: an on-site Local Coordinator, Circuit Riders who rotate across multiple locations on a schedule, Virtual Specialists delivering remote expertise, and a Regional Support Team providing backup.
Developmental Pipeline	A five-stage progression from member to staff member: Service Participant, Occasional Contributor, Regular Volunteer, Staff Apprentice, and Staff Member (compensated).

The source gives no headcount ratio anywhere in this material — no staff-per-member figure and no staff-per-square-footage figure for any of the four models.

The competency framework has three layers. Six core domains — Financial Knowledge, Community Connection, Operational Capability, Educational Facilitation, Technological Facility, and Collaborative Practice — are each assessed across five levels: Emerging, Applied, Proficient, Advanced, and Mastery. Teams are expected to combine a mix of levels across the six domains rather than have every member reach Mastery in each; the source frames this explicitly as a correction adopted after internal review found an earlier version of the framework implicitly required unrealistic universal expertise. Beyond the six core domains, the source names five specialized competencies — Cultural Facilitation, Conflict Resolution, Specialized Financial Expertise, Program Development, and Network Weaving. Each competency has a stated developmental pathway (entry points, learning resources, practice opportunities, an assessment approach, guidance on advanced application), but the source names no specific curriculum or certification requirement for any of them.

Hiring follows three linked practices: **community-integrated recruitment** (relationship-based identification of candidates, community nomination, structured pathways for members to gradually take on larger roles before being formally brought on), **collaborative selection** (current team members and community

members both participate; selection draws on multiple kinds of interaction, including informal observation in a community setting, rather than a single formal interview), and **mutual assessment** (selection treated as two-directional — organizational needs and the candidate’s own goals are both made explicit, and developmental feedback is given regardless of outcome).

Team development operates at three scales: individual (personalized learning plans, mentorship, structured reflection, stretch opportunities), team-level (shared learning processes, conflict-transformation capability, collective adaptation), and cross-team or network-wide (communities of practice across similar roles, staff learning exchanges between locations, shared problem-solving on common challenges).

Compensation rests on a living-wage floor for all contributors, a transparent framework for setting compensation, and regional cost-of-living adjustment. The source also states a narrower gap between the highest- and lowest-paid roles than is typical for the sector, but gives no specific ratio or number for it. Contribution is recognized through paid staff roles, stipended positions, honoraria for specialized expertise, and unspecified non-monetary exchanges; benefits cover financial security, schedule flexibility, access to learning and development, and support for community-engagement time.

Three implementation challenges recur in the source: staffing rural or underserved areas (a smaller local talent pool, gaps in specialized technical expertise, resource constraints on competitive pay, geographic isolation from broader support networks); a structural tension between role clarity and adaptive capacity, and between network-wide consistency and local contextual adaptation; and long-term sustainability (career-pathway design, burnout and workload management, knowledge continuity through staff transitions, and how the staffing model scales as the network grows).

Effectiveness is assessed across three categories — team wellbeing (engagement, sustainability, growth, relationship health), service impact (member feedback, coverage, quality, responsiveness), and community integration (breadth of relationships, trust indicators, how community knowledge flows into the work) — with no numeric target given for any of them.

D. Multi-Purpose Usage Guidance

Running one space across multiple uses is justified on four grounds: **asset optimization**, **community integration**, **service enhancement**, and **financial sustainability**. Asset optimization rests on the observation that a conventional branch stands empty for most of the week. The source offers one figure for this and it is not reproduced here: the number appeared only inside an invented first-person quotation, with no citation and no calculation method, and a figure whose sole provenance is fabricated testimony is fabricated data rather than an unchecked claim. The unverified marker used elsewhere in this section signals a real claim nobody has checked; it would misdescribe this one. The qualitative point stands on its own and does not depend on the number. Community integration treats non-financial visits as building familiarity ahead of financial engagement; service enhancement frames non-financial engagement as a path toward financial-service use; financial sustainability points to shared infrastructure costs and efficiency gains from higher utilization.

Four primary usage categories:

Category	Examples
Financial Service Provision	Individual services and transactions, group financial-education programs, self-service technology access, space for partner organizations.
Community Gathering	Local-group meetings, public forums, recognition events, informal spontaneous gathering.
Learning and Development	Financial-education workshops, general skill-development training, youth programming, peer-to-peer knowledge exchange.
Economic Activity	Showcase space for local entrepreneurs, periodic marketplace events, collaborative workspace, small-business development support.

Scheduling and coexistence rest on four mechanisms: **usage frameworks** (stated priority guidelines and a process for assigning specific activities to specific spaces, with no specific priority order and no numeric balance target given across use types), **coordination systems** (shared scheduling tools, defined transition protocols between uses, a defined approach to resolving scheduling conflicts), **community voice in scheduling** (usage advisory groups of community members, a proposal process for new uses, regular usage reviews), and **partner-use agreements** (formal usage agreements, joint planning, and cost-sharing with regular non-MTU users of the space).

Four categories of physical-design feature enable multi-use: flexible elements (movable furniture, modular components, adjustable lighting and acoustics), storage (activity-specific, organized for rapid setup and takedown between uses), technical infrastructure (systems usable across purposes, with security calibrated to context rather than uniformly maximal), and environmental adaptability (lighting, acoustics, temperature, and capacity flexibility for groups of varying size).

Four usage policy areas recur: accessibility (inclusive physical design, economic access, cultural accessibility, workable scheduling), shared responsibility (explicit use agreements, space-stewardship practices, a defined problem-solving process), safety and security (risk assessment per usage scenario, with a stated principle that security should not create unnecessary access barriers), and resource allocation (cost distribution across uses, equipment-access rules, consumable-supply management).

Coexistence raises four further considerations: regulatory navigation (delineating which activities are subject to which regulation, with no specific regulatory citation given in this material — jurisdictional detail is out of scope for this section); balancing non-financial use against the overall mission; community leadership over usage decisions; and concurrent-use management. On the last point, the source frames its hardest test case — a financial-counseling session, a youth program, and a member needing a quiet space, all running at once — as an illustrative scenario, not a documented incident.

Four illustrative usage archetypes, presented as generic models rather than case studies:

Archetype	Emphasis
Community Financial Center	Individual services, financial-education workshops, budgeting and planning support, financial-technology access.
Community Hub Model	Meeting space, community information exchange, social-connection opportunities, general access to community resources.
Economic Empowerment Center	Small-business and entrepreneur support, workforce-development and job-search assistance, cooperative-development support. The label is the source’s own; it is retained here as a defined term rather than reworded.
Integrated Service Center	MTU financial services co-located with social-service access, basic health information, and educational services.

Usage expands in phases: assessment and planning, core financial functions established first, gradual expansion of additional uses as capacity develops, pilot-testing of new uses, and adaptive development based on experience, with partnership development running alongside. Evaluation covers four categories — utilization metrics (usage hours, user counts, activity diversity, resource efficiency), community-impact assessment (who can access the space, relationship formation, unplanned benefit), financial-sustainability analysis (cost-sharing effectiveness, revenue generated, long-term viability), and mission-alignment evaluation (service quality, member experience, trust development) — and the source gives no numeric target for any of them.

E. Physical-Digital Touchpoints

Physical-digital integration rests on four principles: **human-centered technology** (meant to strengthen human relationships and augment physical interaction rather than replace it, with members retaining control over their own level of engagement), **contextual appropriateness** (technology choices account for a community’s existing access, cultural attitudes, and comfort level, matching the tool to the function needed rather than deploying an advanced one where a simpler tool suffices), **seamless transitions across channels** (a consistent experience whether a member engages physically or digitally, with relationship context and appropriate information flowing between the two), and **progressive introduction** (new technology introduced in graduated steps, opt-in rather than mandatory, with supported exploration for less-comfortable members).

Five concrete integration-point categories:

Category	Function
Member identity and recognition	A unified identity recognized across digital and physical touchpoints; access to relationship context in any channel; authentication designed to avoid friction.
Information access and management	Consistent information across channels; explicit tools for moving information between physical and digital domains; collaborative digital workspaces.
Transaction facilitation	Transactions completable through a member’s preferred channel, including starting in one channel and finishing in another; tools for converting between paper and digital documents.
Learning and capability building	Blended digital and physical learning materials; tools for building digital-skill capability; just-in-time support at the point of need.
Community connection and coordination	Event-management tools for in-person gatherings; communication platforms for the periods between gatherings; tools for visualizing community and network connections.

These are carried by four physical mechanisms. Technology access points — member-facing access stations, staff-support technology, and environmental displays that surface digital information in the physical space — overlap directly with the projection surfaces and status displays §13.4C already describes; this section treats them as the same category of feature, not a second one. Physical-digital interfaces include interactive surfaces, environmental sensors, and physical objects with embedded digital capability; augmented-reality overlays are named as a category but not described as deployed (see Future Directions, below). Technology support zones provide dedicated space for building technology skills, staffed helper stations, peer-support areas, and private space for sensitive digital activity. Infrastructure covers reliable connectivity, accessible power, and environmental controls suited to housing technology.

Four supporting roles bridge the channels: Technology Guides (staff who help members use digital tools), Integration Specialists (make channels interoperate), Community Tech Stewards (peer members who help other members), and Experience Designers (own the end-to-end journey across channels). These are distinct from the consent-terminal and kiosk interactions §13.4C describes; they staff the connective layer beneath those interfaces rather than the interfaces themselves.

Three named implementation challenges: digital divides (access disparities, skill-level variation, generational differences in technology comfort), technology evolution (integration complexity across systems, legacy-technology constraints, future-proofing requirements), and resource constraints (implementation cost, need for specialized expertise, ongoing maintenance burden).

Evaluation covers experience quality (cross-channel journey smoothness, friction points, satisfaction, issue-resolution efficiency), inclusion (access equity, utilization patterns, barriers, adaptability to diverse needs), relationship impact (effect on connection, trust, and depth of engagement), and operational effectiveness (resource-use efficiency, cross-channel information flow, system reliability) — again with no numeric target for any of them.

The source names four directions as exploratory or aspirational, not as deployed capability: ambient computing (technology that recedes into the background, context-aware systems), enhanced reality (informational overlays on the physical environment, shared visualization tools), distributed presence (richer embodied remote participation, ambient awareness of others), and ethical integration (transparency by design, technology framed as preserving rather than undermining member agency). None of the four is described as implemented anywhere in the material.

Conclusion of 13.9

Read together with §13.4C, this section states the physical layer at two grains: §13.4C describes the sensory and consent-interface texture of a single generic hub, and this section states how that hub gets sited, staffed, scheduled across competing uses, and connected to the MTU’s digital systems beneath the kiosk layer. Three gaps recur often enough to name once more, rather than only at each occurrence above: the source gives no floor-area or square-footage figures anywhere in the material, no staff-to-member or staff-to-space headcount ratio for any staffing model, and no numeric scoring or weighting scheme for location selection. Filling any of the three with an estimate would misrepresent the source; this section states that they are unaddressed and leaves them for a later pass with access to a primary source or an operating location.

This section also does not cover community-integration strategy — the source’s parallel material on embedding an MTU location within a community’s existing organizations and relationships. That material sits outside the space, siting, staffing, and technology scope set for this section, and the underlying source file is itself truncated partway through its final subsection. That gap remains open for later work.

§13.10 Jurisdictional Compliance

§13.7 Section C (“Legal and Regulatory Anchoring”) describes the Member Trust Union (MTU) — the network’s credit-union-like financial institution — as operating through a three-layer federated legal wrapper: a Local MTU structured as a “Co-op, Mutual Society, or CDFI-equivalent,” a Regional Network as a “Multi-Stakeholder Cooperative or Umbrella Org,” and a Global Coordination Layer as a “Public Benefit Corporation / DAO LLC.” These are category names, not jurisdiction-specific entity types. §13.7 Section D (“Regulatory Interface Models”) names four compliance domains by function rather than by law — AML/KYC equivalents addressed through verifiable-credential identity and trust-path risk scoring rather than Bank Secrecy Act terminology, tax compliance through flow accounting and smart trust contracts, consumer protection through member-triggered Collective Protective Protocols, and financial-instrument classification through token design intended to avoid securities status. Neither section names a statute, a regulator, or a dollar figure. Chapter 08’s overview (§8, “Legal and Regulatory Architecture”) frames a modular legal architecture for HAOs (Humanized Autonomous Organizations — the network’s coordinating framework) and their constituent units at the same level of abstraction, and §8.1 (“Modular Legal Forms and Multi-Level Structures”) supplies a generic entity-form menu — LLC, worker cooperative, B-Corporation, UK Community Interest Company, and DAO-LLC for United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people; Swiss Verein, Dutch Stichting, nonprofit LLC, federated DAO LLC, and holding cooperative for the HAO Core — again without naming a specific statute, regulator, or figure.

This section supplies the concrete layer beneath that framing, for the MTU specifically: the United States federal and state statutes and regulators a staged source names, a four-state money-transmitter comparison, and European Union, United Kingdom, and Asia-Pacific regulatory material. It does not replace §13.7 or chapter 08’s abstract treatment; it names candidate concrete instances of it. §13.11 continues with the legal vehicle options and phased entity strategy the same source describes.

Source and marker convention. This section is drawn from a single staged extraction of source material located at `~/code/mtu-to`, unreviewed single-pass generative output produced over a year ago and never checked against a primary source — a statute’s own text, a regulator’s own publication, or a state financial-regulator filing. Unlike the source underlying §13.9, the staging extraction reports that none of the six source files reviewed for this material contain invented first-person testimony; nothing is withheld here on that basis. Every named statute, named regulator, dollar figure, bond amount, fee, percentage, and jurisdictional threshold below is nonetheless a real claim nobody has checked, and each carries the marker **(unverified)** once, at its point of introduction — immediately after a named statute or regulator, or via a Status column where the material is tabular. A figure or requirement stated in the same clause as an already-marked statute or regulator is covered by that marker rather than tagged a second time. The marker does not mean a claim is false; it means the claim has not been checked against a primary source, and should be before use. Generic descriptive statements that name no specific statute or regulator and give no figure — for example, that state banking departments generally regulate state-chartered institutions — are not individually marked, consistent with this being a claim about category, not a specific unverified fact.

A. United States — Federal Statutes and Regulators

AML and sanctions. The source names the **Bank Secrecy Act (BSA) (unverified)** as the parent federal framework, requiring a Customer Identification Program (CIP), Customer Due Diligence (CDD) and Enhanced Due Diligence (EDD), Suspicious Activity Reporting (SAR) filing, and Currency Transaction Reports (CTR) for transactions over \$10,000. The **Financial Crimes Enforcement Network (FinCEN) (unverified)** is named as the federal body administering BSA regulation and enforcement. The **Office of Foreign Assets Control (OFAC) (unverified)** is named separately, covering screening against the Specially Designated Nationals (SDN) list, sanctions-program compliance, transaction monitoring, and blocking procedures for matches; the source attaches no figure to OFAC compliance.

Consumer protection. Three federal statutes are named, each with an implementing regulation: the **Truth in Lending Act (TILA) (unverified)** and Regulation Z (unverified) (credit-term disclosure, APR calculation and disclosure, right of rescission, advertising restrictions); the **Electronic Fund Transfer Act (EFTA) (unverified)** and Regulation E (unverified) (electronic-transfer disclosures, error-resolution procedures, limits on consumer liability, preauthorized-transfer requirements); and the **Equal Credit Opportunity Act (ECOA) (unverified)** and Regulation B (unverified) (prohibition on discriminatory lending, adverse-action notice requirements, record retention, fair-lending monitoring). The source attaches no dollar figure or threshold to any of the three.

Federal reporting. Beyond BSA/CTR/SAR, the source names an **NCUA Form 5300 (“Call Report”) (unverified)** as a required financial report for credit-union-structured MTUs; **Home Mortgage Disclosure Act (HMDA) (unverified)** reporting, qualified in the source only as “if applicable”; **Community Reinvestment Act (CRA) (unverified)** documentation, qualified only as “for certain structures”; and a **Foreign Bank and Financial Accounts Report (FBAR) (unverified)** filing, named as an example of foreign-account compliance for cross-border operation. Beyond these named items, the source separately catalogs federal reporting content in more granular form, without tying any of it to a further named report or statute: annual audited financial statements, a BSA compliance certification, business-continuity and disaster-recovery testing results, information-security program assessments, and material outsourcing-arrangement evaluations. Nothing comparable — a specific report name, a specific filing cadence — appears in §13.7, which treats transparency only as a cultural and architectural principle (open ledger views, participatory budgeting) rather than as a set of regulatory filing obligations.

Federal regulatory bodies. Four bodies are named. The **National Credit Union Administration (NCUA) (unverified)** is described as the primary regulator for federally chartered credit unions and as administrator of the **National Credit Union Share Insurance Fund (NCUSIF) (unverified)** — named elsewhere in the source as the applicable deposit-insurance mechanism for a credit-union-structured MTU, with unnamed “private deposit insurance alternatives” mentioned as a fallback for other structures and no specific insurer or coverage limit given. The **Consumer Financial Protection Bureau (CFPB) (unverified)** is described as the federal consumer financial law enforcer. FinCEN is described again in its AML-administration role. The **Federal Reserve System (unverified)** is described as relevant if an MTU interfaces with traditional banking, and for electronic-funds-transfer regulation. That second attribution looks stale rather than wrong: rulemaking authority for the electronic-funds-transfer rules moved to the Consumer Financial Protection Bureau under the 2010 Dodd-Frank Act, which predates the source. It is recorded as the source states it, with the discrepancy noted.

Entity-classification ambiguity. The source frames the MTU model as falling between three existing regulatory categories, with an open question attached to each: a credit union (similar member-ownership structure, a differing loan-distribution mechanism, open questions about field-of-membership requirements); a

peer-to-peer lending platform (shared direct member-to-member lending, a differing risk-pooling approach, regulatory treatment described as varying by jurisdiction); and a mutual aid society (historical precedent for community-based financial support, which the source states modern regulatory frameworks “often fail to account for,” with potential for exemptions or special classification in some jurisdictions). This is a classification argument rather than a statute, a regulator, a figure, or a bond amount, and carries a correspondingly lower verification burden than the material above, so it is not tagged with the marker; it is retained because no equivalent discussion of where the MTU sits relative to existing regulatory categories appears anywhere else in the corpus.

B. United States — State-Level Requirements

Four states appear in the source with money-transmitter figures attached; no other state is covered. Every figure in the table below is **(unverified)**.

State	License Type	Bond Requirement	Application Fee	Annual Renewal	Status
California	Money Transmitter	\$250,000–\$7,000,000	\$5,000	\$2,500	(unverified)
New York	BitLicense / Money Transmitter	\$500,000	\$5,000	\$5,000	(unverified)
Texas	Money Service Business	\$300,000–\$2,000,000	\$2,500	\$1,500	(unverified)
Wyoming	Special Purpose Depository Institution	Risk-based; no figure given	\$15,000	Variable; no figure given	(unverified)

The source states these figures as current “as of” an unspecified date, with no citation to a state statute or regulator publication for any row. It describes state banking departments generically as the primary regulators for state-chartered financial institutions, with requirements it says vary significantly across states, and describes money-transmitter licensing as required in most U.S. states for entities transferring funds between parties. Beyond the four figures above, state-level MTU reporting is described only categorically — annual license renewals, quarterly activity reports, change-in-control notifications, and agent-location reporting — with no specific state statute named for any of the four.

Two further categories of state law are named: state **Unfair, Deceptive, or Abusive Acts or Practices (UDAAP)** laws **(unverified)**, described as often more stringent than federal requirements, with a private right of action in many states and potential for “significant penalties”; and state privacy statutes — the **California Consumer Privacy Act (CCPA)** and **California Privacy Rights Act (CPRA)** **(unverified)**, the **Virginia Consumer Data Protection Act (VCDPA)** **(unverified)**, and the **Colorado Privacy Act (CPA)** **(unverified)** — described only by function (disclosure, consent, and data-subject-rights requirements that vary by state), with no penalty figure given for any of the three.

C. European Union

- **Payment Services Directive 2 (PSD2)** **(unverified)** — Strong Customer Authentication (SCA) requirements, Access to Account (XS2A) provisions, requirements for Payment Initiation Service Providers (PISPs), and Open Banking interface standards.
- **Electronic Money Directive 2 (EMD2)** **(unverified)** — requirements for e-money issuers, safeguarding of customer funds, redemption requirements, and limitations on interest.
- **General Data Protection Regulation (GDPR)** **(unverified)** — legal basis for processing personal data; data-subject rights of access, rectification, erasure, and portability; Data Protection Impact Assessment requirements; breach-notification obligations; and potential fines “of up to 4% of global annual revenue or €20 million, whichever is higher.”
- **5th Anti-Money Laundering Directive (AMLD5)** **(unverified)** — customer due-diligence requirements, beneficial-ownership identification, enhanced measures for high-risk transactions, and requirements for virtual asset service providers.

The source attaches no fee, bond, or capital figure to the EU material beyond the GDPR penalty figure above.

D. United Kingdom

- **Payment Services Regulations (PSR)** **(unverified)** — described as similar to PSD2 with UK-specific variation, a Financial Conduct Authority (FCA) authorization process, safeguarding requirements, and conduct-of-business rules.
- **UK GDPR and Data Protection Act 2018** **(unverified)** — the UK’s post-Brexit GDPR analogue, described as “largely aligned” with EU GDPR with potential for future divergence; guidance issued through the **Information Commissioner’s Office (ICO)** **(unverified)**.
- **Competition and Markets Authority (CMA)** **(unverified)** Open Banking requirements — the UK Open Banking Standard, a directory of participants, technical API specifications, and security profiles and standards.

The source attaches no fee or bond figure to any UK item.

E. Asia-Pacific

- **Singapore — Payment Services Act** **(unverified)** — an activity-based licensing framework, a risk-focused regulatory approach, requirements for digital payment token services, and customer-money protection requirements.
- **Singapore — Personal Data Protection Act (PDPA)** **(unverified)** — consent requirements for data collection and use, purpose-limitation principles, breach-notification requirements, and data-transfer restrictions.
- **Australia — Australian Financial Services License (AFSL)** **(unverified)** — licensing for financial product advice and dealing, compliance with financial-services law, risk-management systems, and an “adequate resources” requirement.

- **Australia — Consumer Data Right (CDR) (unverified)** — Australia’s open-banking framework: data-sharing obligations, consumer consent management, and accreditation requirements.

The source adds only high-level characterization beyond these four items, with no statute or regulator named: Singapore is described as taking a “progressive regulatory sandbox approach”; China is described as exercising “tight control over financial services innovation,” with no statute or regulator named for that characterization; and Australia’s open-banking initiatives are described as similar to the EU approach. No jurisdiction outside the United States, the European Union, the United Kingdom, Singapore, and Australia appears anywhere in the source. China is named only in this one passing characterization, with no regime cited.

Conclusion of 13.10

Read together with §13.7 and chapter 08, this section and §13.11 supply the layer those chapters leave abstract: named statutes and regulators for the United States, the European Union, the United Kingdom, and two Asia-Pacific jurisdictions, with figures for the United States only, plus a four-state money-transmitter comparison. All of it derives from a single unreviewed source and none of it has been checked against a primary source; every named statute, named regulator, dollar figure, bond amount, fee, percentage, and jurisdictional threshold above carries the **(unverified)** marker for exactly that reason, and should be confirmed against the statute’s own text or the regulator’s own publication before any of it is relied upon. Where the source gives a category without a figure — OFAC compliance, the three federal consumer-protection statutes, the UK and EU material beyond the one GDPR penalty figure — this section says so rather than supplying one. §13.11 continues with the ten legal vehicle options and the three-phase entity strategy the same source describes.

§13.11 Legal Vehicle Options

docs/08-legal/01-modular-legal-forms.md (§8.1) already supplies a generic entity-form menu for HAOs (Humanized Autonomous Organizations — the network’s coordinating framework) and their constituent units: LLC, worker cooperative, B-Corporation, UK Community Interest Company, and DAO-LLC for United Micro Enterprises (UMEs) — small, self-managing venture teams of up to ~15 people; Swiss Verein, Dutch Stichting, nonprofit LLC, federated DAO LLC, and holding cooperative for the HAO Core; and Dynamic Enterprise Agreements (DEAs) — versioned operating agreements replacing fixed bylaws — as the mechanism binding entities together. That menu is cited here rather than reproduced; the list above is orientation only. This section supplies a domain-specific one for the Member Trust Union (MTU) — the network’s credit-union-like financial institution: ten named legal vehicle options, with the trade-offs a staged, unreviewed source states for each, followed by the three-phase entity strategy the same source describes for how a single MTU might move between them over time.

Marker convention. As in §13.10, this section is drawn from ~/code/mtu-to , unreviewed source material staged for this corpus and never checked against a primary source. Marking here is done at the row level rather than per mention: every row of the table carries a Status column marked **(unverified)**, which covers each statute, regulator, and claim named anywhere in that row. A row’s advantages, limitations, and regulatory considerations are the source’s own stated claims for that vehicle, not independently verified law. Statutes named in the prose outside the table carry the marker at their point of introduction.

A. Ten Legal Vehicle Options

The source enumerates exactly ten concrete legal vehicle options, grouped into five categories, and names no eleventh.

Entity Form	Category	Advantages (source’s claims)	Limitations	Regulatory Considerations	Status
Credit Union Structure	Traditional financial institution	Established regulatory framework for member ownership; deposit-insurance access through NCUSIF; public familiarity; tax advantages as a 501(c)(14) organization	Restricted field-of-membership requirements; highly regulated operating environment; limited ability to implement innovative loan-distribution mechanisms; capital requirements described as typically starting at \$300,000+; lengthy application process	Federal chartering through NCUA or state-level chartering; comprehensive regulatory compliance; regular examinations and reporting	(unverified)
Industrial Loan Company (ILC)	Traditional financial institution	Federal Deposit Insurance Corporation (FDIC)-insured deposit products; exempt from Bank Holding Company Act restrictions; potentially greater operational flexibility than traditional banks; access to payment systems	Limited availability — the source states that only Utah and Nevada charter ILCs; high capital requirements described as \$10M+ typically; intense regulatory scrutiny and lengthy application process; ownership restrictions in some cases	State chartering with FDIC oversight; comprehensive banking regulation; significant compliance burden	(unverified)
Cooperative Corporation	Cooperative	Democratic member control aligned with MTU principles; flexible, adaptable structure; established legal framework in most jurisdictions; potential tax advantages depending on structure and activities	May require adaptation for financial-services activities; varies significantly by state; may face challenges interfacing with traditional financial systems; often lacks the regulatory clarity of dedicated financial-institution structures	State-level incorporation requirements; potential additional requirements based on activities; may require multiple licenses for financial services	(unverified)
Limited Cooperative Association (LCA)	Cooperative	Hybrid model allowing outside investor capital while maintaining cooperative principles; flexible capital structure; maintains democratic member governance; available in several states, including Colorado, Utah, and Wisconsin	Not available in all jurisdictions; less familiar to regulators and potential partners; may create tension between investor and member interests; relatively new legal form with limited precedent	State-specific statutory requirements; additional requirements based on activities; careful structuring needed to balance member and investor rights	(unverified)
Money Services Business (MSB)	Alternative financial service	Lower barrier to entry than depository institutions; can facilitate payment services and fund transfers; more flexible operational model; potentially faster time to market	Cannot offer deposit accounts; state-by-state licensing requirements; significant compliance costs across multiple jurisdictions; bonding and capital requirements	FinCEN registration at the federal level; state-by-state money-transmitter licensing; AML/BSA compliance program requirements	(unverified)

Entity Form	Category	Advantages (source's claims)	Limitations	Regulatory Considerations	Status
Benefit Corporation (B-Corp)	Alternative financial service	Legal recognition of social mission alongside profit motive; flexibility in corporate structure; enhanced ability to pursue social impact; described as an increasingly recognized legal form across states	Not specifically designed for financial services; still requires applicable financial-service licenses; does not inherently provide cooperative governance; public-benefit reporting requirements	State incorporation requirements for benefit corporations; additional financial-service licensing based on activities; annual benefit-reporting requirements	(unverified)
Network of Interlinked Entities	Innovative hybrid	Tailored legal structures per function (financing, operations, technology); risk isolation between activities; ability to optimize each entity's structure for its purpose; flexibility across multiple jurisdictions	Governance and coordination complexity; additional administrative overhead; potential regulatory concerns about transparency and oversight; complex tax considerations	Demonstrating appropriate separation between entities; clear cross-network governance frameworks; transfer-pricing and inter-entity transaction documentation; potential for consolidated supervision	(unverified)
Platform Cooperative	Innovative hybrid	Digital-first cooperative structure for platform-based services; member ownership of the technology platform; alignment with MTU digital-infrastructure needs; described as a growing ecosystem of supportive resources and communities	Evolving legal framework, still developing in many jurisdictions; may require adaptation of existing cooperative statutes; potential challenges with traditional financing; balancing technology development with member governance	Cooperative incorporation requirements; digital service-specific regulations; data-governance and privacy requirements; cross-border operational considerations	(unverified)
Decentralized Autonomous Organization (DAO)	Decentralized	Programmatic governance and automated operations; potential for direct democratic decision-making; reduced administrative overhead via smart contracts; global accessibility from inception	Uncertain legal status in many jurisdictions; regulatory ambiguity; technical complexity and security considerations; limited precedent for financial-services applications	Evolving regulatory frameworks for DAOs, citing Wyoming's DAO LLC law by name; securities-law implications of tokenization; AML/KYC requirements for decentralized systems; jurisdictional questions regarding applicable law	(unverified)
Mutual Aid Network	Decentralized	Historical precedent for community-based financial support; described as often subject to less restrictive regulation; strong alignment with cooperative principles; focus on member wellbeing over profit	Limited ability to scale while maintaining community connection; may lack legal recognition in some jurisdictions; often unable to offer comprehensive financial services; may face skepticism from traditional financial partners	Varies significantly by jurisdiction; potential religious or fraternal-organization exemptions; limitations on scope of activities; insurance regulation if providing certain benefits	(unverified)

Named regulators and statutes appearing within the table (NCUA, NCUSIF, the 501(c)(14) designation, the Bank Holding Company Act, FDIC, FinCEN, BSA, and Wyoming's DAO LLC law) carry the same **(unverified)** status as the dollar figures and jurisdiction lists in the same rows; the Status column marks this at the row level rather than repeating the marker after each cell-internal mention.

B. Three-Phase Entity Strategy

Beyond the ten-vehicle menu, the source lays out a three-phase implementation strategy for how a single MTU might move between vehicles over time:

- **Phase 1 — Foundation Structure:** a Benefit Corporation with cooperative-like bylaws; Money Service Business licensing for core payment functions; a “clearly defined path” to a more comprehensive structure.
- **Phase 2 — Operational Maturity:** evaluation of credit-union conversion; consideration of a multi-entity structure for different functions; assessment of regulatory-sandbox participation opportunities.
- **Phase 3 — Network Expansion:** development of template legal structures adaptable to different jurisdictions; creation of a network governance framework; standardized agreements for inter-MTU collaboration.

This phased strategy operates on a different axis than §13.7 Section C's three-layer model (Local MTU / Regional Network / Global Coordination Layer). The source's phases are a temporal sequence for a single MTU's entity evolution — Benefit Corporation plus MSB licensing, then evaluation of credit-union conversion, then a templated network — while §13.7's layers are a simultaneous federated structure, with Local, Regional, and Global entities existing at once rather than in sequence. The two are not contradictory, but they are not the same model, and this section does not merge them: whether the phased sequence should be read as a maturation path underneath the existing layer model, or presented as a separate dimension (time versus federation level), is left open for later editorial work.

Conclusion of 13.11

Together with §13.10, this section completes the concrete legal layer beneath §13.7 and chapter 08's abstract treatment, for the MTU domain specifically: ten named vehicle options with source-stated trade-offs, and a phased strategy for moving among them. As throughout §13.10, every named statute, named regulator, dollar figure, and jurisdiction list above is a real claim from a single unreviewed source, not an independently verified one, and carries the **(unverified)** marker for that reason. The qualitative trade-off claims in the vehicle table — that a structure is “flexible,” that a legal form is “evolving,” that regulatory treatment “varies by jurisdiction” — are likewise the source's own assessments rather than verified findings; they are covered by the same row-level Status marker rather than tagged individually, consistent with the convention stated at the top of this section and of §13.10.

15.0 • Advanced

§15.1 Cross-HAO Interoperability and Network Federation

As the HAO model (the network's coordinating framework) matures, its viability no longer depends solely on the integrity of a single network such as the ICN (the reference cooperative business network) or MTU (the network's credit-union-like financial institution). The challenge becomes enabling multiple HAOs, potentially operating in different sectors, regions, or domains, to coordinate and federate without losing their autonomy. Section 15.1 examines how HAO networks can interoperate, how shared standards might emerge, and what architectural principles govern these cross-network relationships.

15.1.1 HAO-to-HAO Governance Protocols

Inter-HAO coordination requires meta-governance protocols that facilitate collaboration between two or more autonomous HAOs while maintaining internal sovereignty. These protocols should define:

- **Recognition Mechanisms:** A formal way for one HAO to acknowledge another as a trusted, values-aligned peer.
- **Federated Agreements:** Lightweight, revocable agreements outlining shared purpose, decision-making processes, and boundary conditions for cooperation.
- **Crisis Collaboration:** Pre-negotiated response protocols for network-wide disruptions (e.g. financial contagion, reputational attacks, external policy threats).

Meta-governance should preserve the polycentric character of each HAO while offering a soft layer for consensus-building, especially across SEPs (joint ventures between teams) that span multiple HAOs.

15.1.2 Inter-HAO Conflict Mediation and Arbitration

Increased interdependence between HAOs increases the likelihood of friction. Conflict resolution pathways between HAOs can maintain continuity without imposing hierarchy. Models might include:

- **Federated Ombudsperson Circles:** Rotating cross-network mediators empowered to facilitate dialogue.
- **Multi-House Councils:** Temporary, issue-specific working groups composed of delegated representatives from involved HAOs.
- **Recursive Consent Loops:** Iterative deliberation protocols designed to surface deeply rooted misalignments and seek resolution without escalation.

Where DAOs often rely on code-based finality, HAOs use human mediation that accounts for culture, power dynamics, and evolving context.

15.1.3 SEP Coordination Across HAO Boundaries

Strategic Enterprise Partnerships (SEPs) represent a key site of interoperability. Cross-HAO SEPs require:

- **Resource Sharing Frameworks:** Agreements on infrastructure, funding, data, or personnel sharing that don't compromise each HAO's internal constraints.
- **Contribution Accounting Protocols:** Ledger-compatible systems that track input and value creation across organizations (e.g., convergent slices, mutual credit bridges).
- **Reputation Portability:** Systems that allow member-level trust and contributions to be recognized across HAOs without duplicating verification.

This supports shared ventures in logistics, research, manufacturing, or policy advocacy across the federated network.

15.1.4 Identity and Trust Federation

For individuals, teams, and UMEs (small, self-managing venture teams) to operate across HAOs, there must be interoperable identity and trust mechanisms. This includes:

- **Decentralized Identity Standards (DID):** Self-sovereign identity frameworks with revocable credentials, endorsements, and privacy-preserving audits.
- **Cross-HAO Trust Graphs:** Merged or interoperable social graphs that map relationships, contribution histories, and endorsements.
- **Layered Trust Signals:** Distinctions between trust inside a given HAO (e.g., UME-level) versus trust accorded by external networks.

These systems must be opt-in, auditable, and resistant to trust-washing or gamified manipulation.

15.1.5 Interoperability Metrics and Performance Standards

To support alignment and coordination, HAOs can develop shared metrics and system health indicators, such as:

- **Interoperability Scores:** A measure of how well a given HAO can collaborate (protocol compatibility, legal flexibility, data openness).
- **Cross-Network Alignment Index:** A dynamic, multidimensional score tracking value alignment, conflict rate, and trust levels across federated HAOs.
- **Ecosystem Contribution Ledger:** A public, time-stamped record of how each HAO contributes to shared assets, knowledge, or infrastructure.

These metrics support transparency and make it easier to detect when one HAO draws disproportionately on shared infrastructure without contributing in return.

Summary

Cross-HAO interoperability and network federation extend the HAO model beyond isolated deployments to collaboration across organizations, geographies, and sectors. This form of federation is designed to preserve local autonomy and value diversity while supporting coordination and resilience at scale.

This section lays the groundwork for further work on interoperable governance standards, federated infrastructure, and decentralized reputation systems.

§15.2 AI-Augmented Governance and Decision Support

As HAOs (the network’s coordinating framework) scale across sectors, geographies, and member complexity, traditional coordination mechanisms — human deliberation, consensus, democratic voting — may encounter bandwidth limits, delayed responses, or decision fatigue. AI-augmented governance extends human judgment and supports coordination at scale without displacing human agency.

Rather than full automation, as in many DAO designs, HAO models prioritize **collaborative intelligence**, where human judgment remains central and AI functions as a context-sensitive support layer. This section outlines architectural principles, design challenges, and opportunities for integrating AI into HAO governance and decision-making.

15.2.1 Design of Transparent, Value-Aligned AI Systems

To align with the ETHICAL framework, any AI deployed in governance must be:

- **Transparent:** Model inputs, training data, and outputs must be auditable and interpretable by human actors. No black-box decision-making.
- **Value-Aligned:** Systems must encode or reflect HAO principles (empathy, accountability, learning) and adapt over time with human oversight.
- **Non-Authoritative:** AI should advise, not decide. Decisions remain under human control, using tools such as explainability dashboards and dissent logging.
- **Context-Aware:** AI must localize its behavior based on the cultural, economic, and ethical norms of each UME (a small, self-managing venture team) or SEP (a joint venture between teams).

Example tools:

- Language models supporting deliberation summaries
 - Recommendation engines for policy proposals
 - Alignment audits detecting value drift in SEPs or UMEs
-

15.2.2 Collaborative Intelligence Network (CIN) Extensions

The CIN is a systemic layer within the HAO responsible for linking AI systems to human workflows. Extensions to the CIN may include:

- **Deliberation Support Agents:** NLP (natural language processing) tools that summarize debate, detect logical fallacies, or highlight underrepresented viewpoints.
- **Ethical Impact Scanners:** Pre-decision audits flagging potential risks to value alignment or unintended stakeholder harms.
- **Knowledge Graphs:** Dynamically updated semantic maps of network knowledge, stakeholder expertise, and historical context.
- **Collective Memory Modules:** Time-stamped logs of decisions, rationales, and downstream effects, surfaced for future reference or retrospective analysis.

This is intended to make AI participatory rather than purely data-driven, supporting clarity, inclusivity, and memory.

15.2.3 Human-AI Consent and Deliberation Protocols

A core principle of HAO governance is the **consent model**: decision-making that seeks the absence of objection rather than majority rule. In this context, AI can:

- **Simulate Outcomes:** Present forecasted consequences of different proposals to support informed consent (including downstream social effects).
- **Facilitate Consent Rounds:** Track evolving positions and flag emerging consensus or conflict zones.
- **Detect Manipulation:** Identify coercion, bias, or persuasion patterns that may invalidate genuine consent.

Protocols should include:

- **AI-Scoped Roles:** Clear boundaries on what decisions AI can influence, suggest, or monitor.
 - **Override Mechanisms:** Any participant or group should be able to nullify or challenge AI outputs.
 - **Accountability Trails:** Immutable logs showing where AI played a role in proposal development or decision refinement.
-

15.2.4 Ethical ML Pipelines for Monitoring Alignment and Participation

AI systems can help continuously evaluate whether the network is functioning in alignment with its core values. Ethical monitoring may include:

- **Participation Equity Indexing:** Assess whether deliberation or resource allocation skews toward specific members, demographics, or power centers.
- **Sentiment and Trust Modeling:** Detect emerging dissatisfaction, misalignment, or systemic distrust based on communication and engagement patterns.
- **Mission Drift Detection:** Identify discrepancies between stated objectives and actual behavior (e.g., a UME optimizing profit at the expense of collaboration).

Design criteria include:

- **Data Minimalism:** Use the least amount of personal data required to derive useful signals.
 - **Feedback Inclusion:** Always allow members to contest, annotate, or reverse AI-driven insights.
 - **Distributed Computation:** Run ethical models at the edge (locally) when possible to preserve sovereignty.
-

15.2.5 Boundaries of Algorithmic Judgment

Despite its utility, AI cannot replace certain functions within HAO governance:

- **Normative Interpretation:** Human beings remain the interpreters of principles such as empathy, harmony, or integrity.
- **Edge Case Ethics:** Situations involving paradox, systemic harm, or moral uncertainty are routed to human deliberation.
- **Conflict Resolution:** While AI can assist with diagnostics, emotional repair and reconciliation remain human tasks.

Refusal, slow paths, and manual override are design features alongside algorithmic acceleration.

Summary

AI in HAO governance is designed to extend capacity without displacing human agency. Systems that augment rather than automate allow HAOs to scale participation while retaining human oversight.

The Collaborative Intelligence Network (CIN) integrates tools for decision support, alignment monitoring, and participatory augmentation. Ongoing work in this area emphasizes auditability, contextual adaptation, and normative transparency, positioning AI as a supplement to governance rather than a replacement for it.

Future directions may include open-source reference AI modules for HAO use, alignment benchmarks for machine-assisted proposals, and federated learning networks across HAOs that maintain data sovereignty.

§15.3 Trust Algorithms and Verification Frameworks

In Humanized Autonomous Organizations (HAOs) — the network’s coordinating framework — trust is not an abstract virtue; it is treated as infrastructure. It determines access, responsibility, and participation scope across all levels: from UME (a small, self-managing venture team) collaboration to capital flows to public market interfaces.

As networks scale, manual, social trust systems reach limits. Algorithmic trust augmentation can support flexible, layered, and context-aware trust systems while limiting continuous surveillance or unconsented use of participant data.

This section defines the architecture, components, and application spaces for trust algorithms in HAO systems.

15.3.1 Progressive Trust Scoring and Non-Binary Verification

Unlike binary identity systems (verified vs. unverified), HAO-aligned trust systems should support **progressive verification** across time and context. Core elements include:

- **Progressive Onboarding:** Initial trust may grant limited permissions (read-only, observer roles). As participation deepens, access expands.
- **Layered Trust Domains:** Individuals can be highly trusted in one SEP (a joint venture between teams) or UME (e.g., logistics) but be newcomers in another (e.g., finance).
- **Risk-Weighted Roles:** Assignments are mapped to trust thresholds (e.g., voting, capital access, mediation authority) with decaying trust curves over inactivity.

This makes trust earned, contextual, and decayable, avoiding centralization or stagnation.

15.3.2 Identity Without Surveillance: Self-Sovereign ID

Traditional identity systems rely on centralized verification, storing sensitive data. HAOs require:

- **Self-Sovereign Identity (SSI)** models using decentralized identifiers (DIDs) and verifiable credentials
- **Zero-Knowledge Proofs (ZKPs)** for verifying facts without disclosing personal information (e.g., “member has 6+ months contribution history”)
- **Selective Disclosure:** Participants choose which claims to share, when, and to whom
- **Portable Credentials:** Credentials can travel across HAOs, with local overrides and time-bound scopes

Trust can be machine-verifiable and human-governed, maintaining privacy by default while supporting movement across the ecosystem.

15.3.3 Contextual Risk Assessment for Peer-to-Peer Finance

Trust is most operationalized in HAOs when it mediates access to capital, from micro-loans to collective investment. Trust algorithms can inform:

- **Risk Weighting:** Lending terms based on trust tier, prior repayment patterns, and network endorsement
- **Staggered Lending:** Smaller commitments first, followed by scaling investment as trust and history accrue
- **Relational Guarantees:** Loans guaranteed not by collateral but by trust relationships (multi-signed endorsements or shared fallback groups)

This supports capital flows based on embedded reputation and mutual accountability rather than collateral requirements.

15.3.4 Combining Relational and Behavioral Trust Signals

A trust algorithm doesn't rely solely on transactions or endorsements. It integrates multiple classes of data:

- **Relational Trust:** Who vouches for you, how long you've worked together, degree of mutual risk
- **Behavioral Trust:** Timeliness of delivery, adherence to community norms, responsiveness, participation in governance
- **Temporal Signals:** Trust should mature and decay over time based on engagement patterns
- **Situational Adaptation:** Trust thresholds vary by task (e.g., voting vs. budget access)

Trust algorithms must remain interpretable, enabling users to contest or correct misjudgments and biases.

15.3.5 Trust Graphs as Infrastructure (Especially in MTUs)

Trust relationships across members, UMEs, and SEPs form an evolving **trust graph**. In the MTU (the network's credit-union-like financial institution) context especially, this graph becomes a financial substrate:

- **Weighted Edges:** Represent varying degrees of endorsement, verification, or shared history
- **Community Trust Pools:** Groups whose aggregate trust supports access to shared credit
- **Pathfinding:** Can user X reach trusted status with group Y via trusted intermediaries?
- **Anomaly Detection:** Identify trust-farming behavior or coordinated attempts to manipulate the graph for disproportionate gain

Trust graphs support dynamic access control, risk modulation, and distributed permissioning across HAOs.

Summary

Trust algorithms in HAO ecosystems are designed to enhance, distribute, and contextualize human judgment rather than replace it. Treating trust as dynamic, relational, and composable is intended to avoid both rigid credentialism and undifferentiated treatment of participants.

Trust functions as infrastructure supporting peer-based finance, access to roles, and coordination without a central enforcing authority. The design challenge is combining technical rigor with attention to relationships between participants.

Future research should explore:

- **Composable trust libraries** tailored to different HAO domains (e.g., MTU finance vs. ICN logistics, where ICN is the reference cooperative business network)
- **Federated trust architectures** with opt-in cross-HAO bridges
- **Post-quantum ZKP trust frameworks** for long-term resilience
- **Gameable behaviors and failure modes**, including sybil attacks, trust inflation, or social engineering within algorithmic ecosystems

§15.4 Mutual Credit Systems and Internal Economies

HAOs (the network's coordinating framework) aim to decouple value creation from external market dependency by enabling trust-based internal economies. One core mechanism for achieving this is a **mutual credit system**: a closed-loop, ledger-based exchange framework where participants earn and spend based on collective trust rather than fiat liquidity or collateral.

This section outlines how mutual credit can function within and across HAOs, the design requirements to avoid failure modes, and its role in supporting internal economies that don't depend on external capital.

15.4.1 Integration of Mutual Credit with UME Operations

Each UME (a small, self-managing venture team) operates as a node of value creation. A mutual credit system allows UMEs to:

- **Transact without external capital:** Goods, services, and labor can flow based on trust-backed credits instead of cash.
- **Issue credit upon creation:** Rather than pre-funding work, UMEs or members earn credits when value is provided.
- **Operate within local thresholds:** Each UME may set internal spending or issuance caps to prevent overextension.

Example: A logistics UME needs design work. It pays a creative UME in mutual credits, which can later be spent with a tooling UME, creating a circular value flow.

15.4.2 Clearinghouse Protocols for Multi-UME Exchange

In large HAO networks, value must be transferable across UMEs without constant bilateral negotiation. This requires:

- **A Federated Credit Clearinghouse:** A protocol-layer ledger that tracks credits across all participating UMEs and reconciles debits/credits network-wide.
- **Multi-Party Transactions:** A system to support triangular or higher-order exchanges (e.g., X pays Y, Y pays Z, Z repays X).
- **Trust-Weighted Routing:** Credits may move along more trusted paths in the network graph, reducing the need for direct bilateral trust.

This creates liquidity without cash, enabling collaboration at scale without centralized issuance.

15.4.3 Reserve Balancing and Inflation Control

Poorly designed mutual credit systems can suffer from inflation, credit hoarding, or system lock-up. Safeguards include:

- **Credit Limits and Drift Windows:** Entities may go negative up to a threshold and must return to balance within time constraints.
- **Diminishing Credit Utility:** Credits decay in value or utility if hoarded, incentivizing circulation.
- **Backing Pools:** A fraction of fiat revenue from external-facing SEPs (joint ventures between teams) can backstop system confidence.
- **Liquidation Triggers:** If a UME collapses or cannot repay, credits can be absorbed by mutual insurance or redistributed.

This keeps credit anchored in real value and community trust, rather than artificial scarcity or speculative flow.

15.4.4 Market Design for Localized Production

Mutual credit systems can be designed to support local production and exchange. Features include:

- **Incentives for Internal Sourcing:** Higher credit value for purchasing from within the MEE (the network's protected internal economy) or UME network.
- **Buy-Local Weighting:** Internal pricing mechanisms that adjust based on distance, production method, or UME classification.
- **Credit-Only Marketplaces:** Internal platforms where UMEs and members can buy/sell services exclusively in mutual credit, increasing internal liquidity.

The system weights allocation toward needs and trust over profit, supporting production and resilience in supply chains.

15.4.5 Regulatory Navigation and Legal Considerations

Operating internal currencies and credit systems raises compliance questions. Design considerations include:

- **Legal Distinctions:** Mutual credit is not debt, interest-bearing, or speculative, which may exempt it from financial regulation in many jurisdictions.
- **Closed-Loop Framing:** As long as credits can't be converted to fiat directly, they may be considered as reputation points or internal vouchers.
- **Taxation Models:** HAOs must account for how mutual credit transactions impact tax liabilities (e.g., if credits are tied to real services).
- **Governance Layer:** Credit issuance, redemption, and disputes are governed by consent-based HAO protocols, not centralized authorities.

Careful architectural, legal, and semantic design helps mutual credit systems operate within ambiguous regulatory territory, while preparing pathways for future legal clarity.

Summary

Mutual credit systems give HAOs a financial substrate where credit issued corresponds to value backed by goods and services. Paired with trust graphs and federation protocols, this supports scaling without dependence on debt or outside investor control.

The internal economy runs on different mechanisms than fiat currency: it links social trust and collective coordination to how value moves between participants.

Future extensions include:

- Cross-HAO credit clearing protocols
- Trust-weighted credit issuance algorithms
- Reputation-linked credit staking models
- Commons-linked credit pools for infrastructure investment

Together, these systems reduce HAOs' dependency on external markets, building an internal economy based on trust, contribution, and mutual resilience.

§15.5 Bioregional and Sector-Specific HAO Implementations

While the ICN (the reference cooperative business network) and MTU (the network's credit-union-like financial institution) represent generalized blueprints for economic and financial coordination, HAOs (the network's coordinating framework) can also adapt to context. Whether the focus is a bioregional economy, a public health system, or a knowledge commons, the HAO pattern can be adapted to different domains without changing its fundamental architecture.

This section outlines the structural adaptations, cultural considerations, and technical mechanisms for deploying HAOs in regions, industries, or public service domains beyond enterprise and finance.

15.5.1 HAO Design for Agricultural and Resource Systems

Agricultural and ecological stewardship are contexts suited to HAO implementation, given:

- Embedded local knowledge
- Long-term resource dependencies
- Need for polycentric governance

Key adaptations include:

- **Land Stewardship UMEs** (small, self-managing venture teams): Entities focused on farming, watershed protection, or forest maintenance, governed by community-based SEPs (joint ventures between teams).
- **Commons-Based Coordination**: Shared seed banks, water rights, or equipment pools managed through dynamic agreements and mutual credit systems.
- **Participatory Monitoring**: Collaborative Intelligence Network (CIN) extensions for reporting soil health, climate patterns, and crop yields, with AI-assisted alerts for climate shocks or overuse.

This supports region-specific governance over vital resources, grounded in ecological feedback and intergenerational ethics.

15.5.2 Regional Commons Governance and Infrastructure

HAOs can govern shared public infrastructure at the bioregional or city scale by adapting:

- **Nested Representation Models**: Geographic or functional sub-HAOs (e.g., transportation, housing, waste) with overlapping stakeholder groups.
- **Asset Stewardship SEPs**: Multi-stakeholder ventures that manage infrastructure through consent-based governance and reinvestment mechanisms.
- **Local Credit Systems**: Regionally bounded mutual credit ecosystems that encourage reinvestment in community-scale needs.

Example: A public transportation HAO governed by riders, workers, municipalities, and local manufacturers, aligned via a shared DEA (a versioned operating agreement replacing fixed bylaws) and revenue reinvestment loop.

15.5.3 Public Sector and Service-Oriented HAOs

Public services, including education, health, and civic data, can adopt HAO patterns organized around participatory governance:

- **Service UMEs**: Schools, clinics, or research institutions operating semi-autonomously under shared ethical and performance agreements.

- **Accountability Protocols:** VAM (ongoing checks that actions match stated principles) systems that keep services accountable to accessibility and equity metrics.
- **Consent-Based Oversight Councils:** Replacing top-down bureaucracies with multi-perspective review boards representing both providers and recipients.

HAOs offer a structural alternative to privatization and state centralization, keeping public systems responsive without extracting profit from service delivery.

15.5.4 Knowledge, Arts, and Cultural Sector HAOs

The knowledge commons is a fit for HAO implementation given its need for:

- **Distributed contribution and ownership**
- **Contextual valuation**
- **Monetization shared across contributors**

Model features include:

- **Creative SEPs:** Multi-disciplinary teams producing open-access media, publications, or educational tools under collective revenue-sharing models.
- **Credited Contribution Systems:** Tracking of citations, code commits, or instructional impact within a trust-based ledger for long-term equity accrual.
- **Commons Finance Pools:** MTU-aligned credit systems funding creative work through collective prioritization and patronage systems.

This lets knowledge work continue without depending on platforms that capture most of the value created, or on IP lock-in.

15.5.5 Integration with Indigenous and Traditional Governance

HAOs are not intended to override existing systems of cultural self-governance. In many cases, the HAO can function as a bridge or interpreter layer, offering:

- **Consent-Respecting Protocol Translation:** Mapping traditional decision-making forms into compatible digital governance structures.
- **Autonomy Preserving Zones:** UMEs or sub-HAOs under indigenous authority, federated via memoranda of mutual recognition, not central control.
- **Land-Linked Governance:** Embedding territory and lineage rights into the HAO structure with non-financial authority tokens or stewarding roles.

This keeps HAO structures compatible with self-determined cultural and ecological governance, rather than replacing it.

Summary

Bioregional and sector-specific HAOs allow the framework to adapt to terrain, tradition, and domain. Whether managing riversheds, running community clinics, or stewarding a local media network, the HAO pattern offers:

- Decentralized authority
- Value-aligned coordination
- Flexible participation and trust systems
- Sovereignty without isolation

These implementations extend the HAO pattern beyond work coordination to broader domains of self-governance.

Future research may focus on:

- Comparative field studies of regional HAO variants
 - Hybrid HAO-public institution models
 - Trust and consent protocols across cultures and governance ontologies
 - Environmental and social impact metrics tailored to bioregional HAOs
-

§15.6 Systems Simulation, Modeling, and Stress Testing

The Humanized Autonomous Organization (HAO) — the network’s coordinating framework — introduces a systems architecture combining polycentric governance, trickle-out economics, trust-anchored finance, and AI-augmented deliberation. This complexity calls for simulation and modeling techniques to validate design assumptions, anticipate failure modes, optimize systemic responses, and support deployment.

Section 15.6 presents a structured approach to modeling HAOs as complex adaptive systems, drawing from systems dynamics, agent-based modeling, computational social science, and scenario-based resilience engineering. This work supports prototyping and de-risking real-world HAO implementations.

15.6.1 Agent-Based Models for UME and SEP Interactions

At the core of HAO simulation is **agent-based modeling (ABM)**, where autonomous agents — UMEs (small, self-managing venture teams), members, and SEPs (joint ventures between teams) — interact within defined rules, producing emergent systemic behavior.

Agent Types:

- **UME Agents:** Defined by lifecycle stage, production type, trust reputation, credit balance, and value alignment score.
- **Member Agents:** Individuals with roles, trust scores, skills, and alignment tendencies.
- **SEP Agents:** Inter-UME collaborations governed by DEA (a versioned operating agreement replacing fixed bylaws) extensions, resource sharing rules, and time-based triggers.
- **HAO Meta-Agent:** Represents coordinating functions, not hierarchical control; performs dynamic redistribution, audits, or triggers systemic alerts.

Simulation Objectives:

- Observe the propagation of trust or misalignment through interconnected agents.
- Detect overcentralization tendencies or idle capital in mature UMEs.
- Identify optimal thresholds for inter-UME collaboration via SEPs.

Tools: Mesa (Python), NetLogo, GAMA Platform, or multi-agent extensions in Rust/Scala for high-scale throughput.

15.6.2 Dynamic Simulation of Outward Allocation and Diminishing Contributions

In the ICN (the reference cooperative business network), the financial logic inverts traditional flows: capital enters at the center (the HAO) and is distributed outward to value creators (UMEs). Separately, diminishing contributions return to shared infrastructure at decreasing rates.

Key Parameters to Simulate:

- Initial capital deployment schedules
- Return-to-center (HAO) percentages per UME maturity phase
- Liquidity curves under delayed UME productivity
- Reinvestment thresholds for surplus UMEs
- SEP-specific cash flow allocation

Stress-Test Variables:

- Over-saturation of early-stage UMEs
- UME collapse and unrecovered investments
- Coordination lags across SEPs with asynchronous financial cycles
- Public market interface volatility affecting the network

This model informs reserve design, credit issuance safety, and capital routing policies.

15.6.3 Failure Mode and Antifragility Scenario Modeling

A central HAO principle is **antifragility**: systems that grow stronger under stress rather than merely surviving it. Simulation includes intentional stress induction and system validation.

Failure Modes to Explore:

- **Trust collapse** in a regional MTU (the network's credit-union-like financial institution) and its ripple effect
- **Governance fragmentation** from DEA versioning conflicts
- **AI model misalignment**, causing decision drift
- **Reputation hoarding or sybil attacks** in trust networks

Antifragility Patterns to Test:

- New UMEs formed by splitting failing ones
- Adaptive distribution rebalancing during SEP failure
- Rapid democratic reconstitution of roles during value misalignment
- Member migration and local reintegration after UME collapse

Evaluation Criteria:

- Time to recovery (TTR)
 - Systemic integrity retention
 - Equitable redistribution of loss and opportunity
-

15.6.4 Human-in-the-Loop Scenario Walkthroughs

Beyond purely computational modeling, HAO simulation must include **human-in-the-loop (HITL)** processes for qualitative and ethical dynamics.

Scenario Design:

- Interactive workshops with real stakeholders roleplaying UME collapse, SEP renegotiation, or AI-flagged value drift.
- Mixed-method walkthroughs combining simulation telemetry with participant emotional responses.
- Use of collaborative foresight tools (e.g., Miro, Kumu, Loomio) for distributed reflection.

Goals:

- Detect misalignments between algorithmic recommendations and human expectations
- Surface latent conflicts or ethical edge cases not present in model logic
- Build cultural fluency and resilience into governance norms before deployment

HITL methods help simulation account for ethical as well as technical considerations.

15.6.5 Game-Theoretic Modeling of Incentives and Strategic Behavior

All systems are vulnerable to strategic manipulation. Simulating incentive alignment under game-theoretic conditions is one method for studying this risk.

Use Cases:

- UME over-reporting value for premature maturity status
- Members “playing the trust graph” to unlock capital access faster
- Public market interfaces attempting to pull governance upward via investment leverage

Approaches:

- Nash equilibrium detection under various trust issuance algorithms
- Iterated games with punishment/reward dynamics based on ETHICAL violations
- Mechanism design simulations to assess resilience of mutual credit architecture

Interventions:

- Adaptive friction: increasing verification friction under high volatility
- Community veto mechanisms for stopping “gaming” at scale
- Penalty decay: system-encoded reduction of privileges over time after abuse

This is meant to make HAOs more robust against incentive misalignment, rather than dependent on idealized behavior.

Summary

Simulation, modeling, and stress testing are intended to surface weaknesses before they appear in deployment. A modeling strategy for HAOs includes:

- Quantitative agents and flows (ABM, trickle-out finance)
- Qualitative reflection (HITL, cultural walkthroughs)
- Systems thinking (feedback loops, risk dynamics)
- Game theory and behavioral prediction

The simulation layer reflects the same principles as the HAO itself: participatory, adaptive, pluralistic, and human-centered. Models aim to make complexity legible and support preparedness, rather than reduce it.

Future work includes:

- A shared HAO simulation stack (open-source and modular)

- Network-wide stress drills across federated HAOs
- Embedded observability layers for real-time mirroring of modeled dynamics
- An open repository of observed HAO failure/recovery cases (similar to postmortems in software SRE)

Appendices

Appendix: Visual Concepts for Section 13.5 – Member Experience and Compensation Systems

This appendix outlines the full suite of visual metaphors and representations to accompany Section 13.5 of the MTU framework. These diagrams aim to enhance conceptual clarity and support downstream implementation.

Figure 1: Member Lifecycle Spiral

Type: Spiral Diagram
Purpose: Visualize non-linear, evolving stages of member engagement.
Structure:

- Central point: **Initiate**
- Outward progression: Participant → Contributor → Steward → Integrator → Elder
- Arcs show lateral transitions, trust repair paths
- Gradient color = trust depth

Metaphor: *Tree ring growth* — Membership maturity expands concentrically with time and trust.

Figure 2: Compensation Architecture Stack

Type: Layered Stack Chart
Purpose: Depict the four interlocking compensation components.
Structure:

- Baseline Livelihood (bottom layer)
- Performance Distributions (next layer)
- Network Profit Sharing (next layer)
- Long-Term Equity (top layer)
- Vertical flow arrows showing member input (labor, governance, care)

Metaphor: *Terraced farmland* — Layered resilience fed by diverse inputs.

Figure 3: Trust-Based Access Wheel

Type: Radial Access Graph
Purpose: Show how deeper trust enables broader system access.
Structure:

- Center = Identity and Trust
- Rings = Access levels: basic services, peer lending, pool governance, equity mechanisms
- Paths radiate based on verified trust graph depth

Metaphor: *Mandala or iris* — Access expands with trust relationships.

Figure 4: Equity Vesting Curve Comparison

Type: Line Graph
Purpose: Compare MTU’s adaptive vesting with corporate linear vesting.
Structure:

- X-axis = Time
- Y-axis = % Equity Ownership
- Two curves:

- Corporate (dotted, fixed 4-year ramp)
- MTU (solid, trust-weighted, with plateau + decline under disengagement)

Metaphor: *River vs. irrigation pipe* — Organic accrual vs. rigid disbursement.

Figure 5: Member Compensation Streams Map

Type: Sankey Diagram

Purpose: Track how diverse member contributions flow into compensation types.

Structure:

- Inputs: Relational Labor, Skills, Governance, Mutual Aid
- Outputs: Livelihood Stipend, Timebank Credits, Equity, Recognition Tokens
- Flow width = contribution weight

Metaphor: *Aqueduct network* — Different channels irrigate specific fields.

Figure 6: Intergenerational Compensation Loop

Type: Circular Loop Diagram

Purpose: Illustrate how compensation flows across generations.

Structure:

- Nodes: Elder → Commons Fund → Dependent Proxy → Long-Term Pool → Elder
- Continuous loop with optional branches for memorial or steward trusts

Metaphor: *Seed vault or heirloom system* — Values and capital circulate through time.

Figure 7: Ethical Filters Overlay

Type: Filter Diagram

Purpose: Show ethical gates applied to any compensation proposal.

Structure:

- Input node = Proposed plan
- Filters:
 - Transparency Audit
 - Deliberative Governance
 - Non-Exploitative Logic
- Output = Validated Compensation Flow

Metaphor: *Lens calibration system* — Ensures alignment and clarity through ethical focus.

Interactive Visual Concepts (Optional Digital Tools)

1. “My Trust Graph” Visualization

- Interactive node-edge graph
- Hover over edges to view relationship trust scores
- Displays which compensation types each trust path enables

2. “Contribution Ledger” Timeline

- Timeline-based ledger
- Color-coded entries for different contribution modes (e.g. blue for governance, green for care work)

3. “Earning Journey Simulator”

- Input sliders for hours/week of different contribution types
- Output dashboard: equity growth, monthly payout, trust score increase

Worker-Owned Café (Portland, OR): UME Viability Models

Model basis:

- 10 worker-owners, all full-time
- Living wage target in Portland: **\$45K–\$60K/year** per worker
- All workers are owners and share profits
- All surplus is reinvested first in maintenance/upgrades, then redistributed

Shared Assumptions:

Parameter	Value
Staff	10 worker-owners
Hours	40/week
Open days	360/year (closed 5 holidays)
Revenue mix	80% food/beverage, 20% event/retail
HAO Network Fee (mature)	10%
Reserve Contribution	5% (equipment, repairs)
Profit redistribution	After infra + reserve

Model A: Modest Success (Surviving, Not Thriving)

Category	Amount
Gross Revenue	\$750,000
COGS (food, drink)	\$180,000 (24%)
Operating Expenses	\$375,000 (50%)
HAO Contribution (10%)	\$75,000
Reserves (5%)	\$37,500
Net Available for Distribution	\$82,500
Per Worker Take-Home (Annual) ~\$53,250 (\$45K salary + \$8.25K surplus)	

Summary: The café pays Portland’s living wage baseline. Slight buffer for emergencies. Surplus is modest, but morale and retention may depend on shared values and stability, not just cash.

Model B: Stable Local Favorite

Category	Amount
Gross Revenue	\$1,050,000
COGS	\$250,000
Operating Expenses	\$450,000
HAO Contribution (10%)	\$105,000
Reserves (5%)	\$52,500
Net Available for Distribution	\$192,500
Per Worker Take-Home (Annual) ~\$64,250 (\$52K salary + \$12.25K surplus)	

Summary: Sustainable and slightly competitive wages. Budget for better equipment or a training program. Can support better retention and social programs (e.g. shared child care).

Model C: High-Performing Community Hub

Category	Amount
Gross Revenue	\$1,400,000
COGS	\$300,000
Operating Expenses	\$500,000
HAO Contribution (10%)	\$140,000

Category	Amount
Reserves (5%)	\$70,000
Net Available for Distribution	\$390,000
Per Worker Take-Home (Annual) ~\$84,000 (\$60K salary + \$24K surplus)	

Summary: Above-average wages in Portland foodservice. Room for margin smoothing, community investment, rotating leadership stipends, or even rotating sabbaticals. Becomes a *destination workplace*.

Notes on Realism & Strategy:

- Portland café revenue per square foot averages **\$300–\$600/year**, so these models assume ~1,500–2,000 sq ft of well-utilized space with seating and some retail
 - Seasonal income variation assumed smoothed by **SEP partnerships**: e.g., pop-ups, catering, or evening events
 - All models maintain **worker equity growth** through retained earnings and reinvestment into tools, spaces, and professional development
 - All models preserve **HAO contributions**, which help fund network-wide support, training, and emergency buffers
-

Updated Models: Customer Volume & Order Value Integration

Model basis: 360 operating days/year.

Model A: Modest Success

Metric	Value
Target Annual Revenue	\$750,000
Average Order Value (AOV)	\$8.50
Required Orders per Day	~245
Seating Capacity	~25–30 (turns ~8–10/day)
Hours Open	10 hrs/day
Orders per Hour (avg)	~25/hour

✓ Operational Notes:

- Modest but steady foot traffic
 - Likely reliant on commuter morning rush + lunch spikes
 - Off-hours likely lean; relies heavily on local regulars
-

Model B: Stable Local Favorite

Metric	Value
Target Annual Revenue	\$1,050,000
Average Order Value (AOV)	\$9.50
Required Orders per Day	~308
Orders per Hour (avg)	~31/hour

✓ Operational Notes:

- Better AOV via food, add-ons, upselling (e.g. local pastries, merch, etc.)
 - May include some catering/SEP activities that sooth out the hourly curve
 - Weekend brunch traffic important here
-

Model C: High-Performing Community Hub

Metric	Value
Target Annual Revenue	\$1,400,000
Average Order Value (AOV)	\$10.50
Required Orders per Day	~370
Orders per Hour (avg)	~37/hour

✓ Operational Notes:

- Requires high-throughput workflow (two baristas, full kitchen, strong point-of-sale)
- Likely includes **evening programming**, SEPs, and possibly alcohol license or event space rental

- Strong loyalty program or subscription/membership model likely in place

Summary Table

Model	AOV	Orders/Day	Orders/Hour	Monthly Revenue
Modest Success	\$8.50	245	~25	\$62,500
Stable Local Favorite	\$9.50	308	~31	\$87,500
High-Performing Community Hub	\$10.50	370	~37	\$116,667

👉 Observations & Strategic Levers

- Every \$1 increase in **AOV** lowers the need for ~35 daily orders
- **SEPs** (e.g. hosting UME events, renting to therapists at night, or art collectives) can **boost revenue without increasing foot traffic**
- Higher AOV may include **subscriptions** (monthly coffee memberships), **retail bundles**, or **add-on services** (e.g., coworking table fees, wellness pop-ups)

Citation Audit

This appendix records the primary-source audit of every work in the corpus's 48-record pre-audit bibliography and four supplemental citations discovered by rebuilding the live citation inventory. The unit of analysis is a citation occurrence: an inline claim, a reference-list placement with a demonstrable scope, or its consolidated-manuscript duplicate. **removed** rows preserve actual deleted or superseded pre-audit attributions; ordinary line-number shifts are not counted as removals.

Scope and method

Metadata was checked against the publication, publisher or journal record, DOI registry, or an official author, university, government, or institutional archive. Claim support was checked against the original publication, an official full-text version, or a publisher-hosted abstract no broader than the audited claim. Search snippets, retailer listings, staged summaries, and generated claim extractions were discovery aids only.

For malformed or uncorroborated identities, the search covered the relevant publisher or journal, DOI and bibliographic registries, author or institutional archives, and likely first-party repositories. **not verifiable** means that boundary produced a genuine identity or access limit; it is not used for an inspected source that failed to support a claim. No nearby source was silently substituted.

Verdicts are **supported** (the source establishes the bounded claim), **partially supported** (only a narrower or corrected form is established), **unsupported** (inspected evidence does not establish the attribution), **general reference** (a broad relationship is established but no claim is attached), and **not verifiable** (the documented primary-source boundary did not expose enough evidence to decide).

Summary

- **Baseline records:** 48
- **Supplemental records:** 4
- **Final canonical works:** 39
- **Live occurrences:** 180
- **Removed occurrences:** 159
- **Replacements:** 0
- **Removals:** 12
- **Unresolved identity or access limits:** 8
- **Supported:** 37
- **Partially supported:** 33
- **Unsupported:** 14
- **General references:** 233
- **Not verifiable:** 22

The five verdict totals include both the current corpus and its preserved pre-audit history. The 180 live occurrences comprise 35 **supported** and 145 **general reference** verdicts; the 159 removed occurrences comprise 2 **supported**, 33 **partially supported**, 14 **unsupported**, 88 **general reference**, and 22 **not verifiable** verdicts.

The final bibliography contains 39 distinct works after twelve removals and the B035-to-B032 merge. Eight removed records retain **not verifiable** findings because no exact primary record was located within the documented boundary. These are unresolved access or identity limits, not live citations.

Evidence ledger

B001 — Money-Go-Rounds: The Importance of Rotating Savings and Credit Associations for Women

- **Canonical citation:** Ardener, S., & Burman, S. (Eds.). (1995). *Money-Go-Rounds: The Importance of Rotating Savings and Credit Associations for Women*. Berg.
- **Metadata status:** verified as an edited 1995 Berg volume; the 1995 catalogue carries the expanded subtitle, while Routledge's current reissue page shortens it to The Importance of ROSCAs for Women
- **Primary record:** 1995 Berg catalogue record (controlling record for the expanded title, editor responsibility, year, and imprint); Routledge current-reissue record (current publisher record with the shortened subtitle)
- **Disposition:** retained
- **Defect tags:** money-go-rounds-editor-role
- **Forbidden live strings:** Ardener, S., & Burman, S. (1995)

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B001-O001	live	docs/13-member-trust-union/02-financial-product-design.md:155	ardener burman 1995	Reference-list placement provides bounded context on rotating savings associations.	general reference	Publisher opening-chapter abstract defines regular contributions with the fund distributed to contributors in turn.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B001-O002	live	docs/13-member-trust-union/02-financial-product-design.md:93	ardener burman 1995	ROSCAs provide a rotating-contribution precedent for the corpus's mutual-aid design.	supported	Publisher opening-chapter abstract defines regular contributions with the fund distributed to contributors in turn.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B001-O003	live	docs/humanized-autonomous-organization.md:5291	ardener burman 1995	ROSCAs provide a rotating-contribution precedent for the corpus's mutual-aid design.	supported	Publisher opening-chapter abstract defines regular contributions with the fund distributed to contributors in turn.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B001-O004	live	docs/humanized-autonomous-organization.md:5355	ardener burman 1995	Reference-list placement provides bounded context on rotating savings associations.	general reference	Publisher opening-chapter abstract defines regular contributions with the fund distributed to contributors in turn.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B001-O005	removed	docs/13-member-trust-union/02-financial-product-design.md:93	ardener burman 1995	The pre-audit sentence attributed ROSCAs, tontines, kibbutz-like redistribution, and the MACI design jointly to Ardener and Burman.	partially supported	The publisher's chapter 1 abstract establishes regular contributions and rotating distribution, but not the other precedents or the corpus's MACI implementation.	Supersede the bundled attribution with the bounded ROSCA precedent and identify the remaining synthesis as corpus-authored.
B001-O006	removed	docs/humanized-autonomous-organization.md:5297	ardener burman 1995	Consolidated pre-audit duplicate of the bundled ROSCA, tontine, kibbutz, and MACI attribution.	partially supported	The publisher's chapter 1 abstract establishes regular contributions and rotating distribution, but not the other precedents or the corpus's MACI implementation.	Supersede the bundled attribution with the bounded ROSCA precedent and identify the remaining synthesis as corpus-authored.
B001-O007	removed	docs/13-member-trust-union/02-financial-product-design.md:155	ardener burman 1995	The pre-audit reference omitted the editors' role and named Berg Publishers rather than the verified Berg imprint.	general reference	The publisher and 1995 catalog records identify Ardener and Burman as editors of the Berg volume.	Supersede the incomplete responsibility and publisher form with the canonical citation.
B001-O008	removed	docs/humanized-autonomous-organization.md:5361	ardener burman 1995	Consolidated pre-audit duplicate omitted the editors' role and used the superseded publisher form.	general reference	The publisher and 1995 catalog records identify Ardener and Burman as editors of the Berg volume.	Supersede the incomplete responsibility and publisher form with the canonical citation.

B002 — Network Society and Future Scenarios for a Collaborative Economy

- **Canonical citation:** Kostakis, V., & Bauwens, M. (2014). *Network Society and Future Scenarios for a Collaborative Economy*. Palgrave Pivot/Palgrave Macmillan. <https://doi.org/10.1057/9781137406897>
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** Springer Nature book and DOI record
- **Disposition:** retained
- **Defect tags:** network-society-author-order
- **Forbidden live strings:** Bauwens, M., & Kostakis, V. (2014)

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B002-O001	live	docs/13-member-trust-union/03-governance-and-risk-distribution.md:151	kostakis bauwens 2014	Reference-list placement provides bounded context on P2P infrastructures and commons scenarios.	general reference	Publisher contents, chapters 3–4, establish P2P infrastructure and commons scenarios, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B002-O002	live	docs/humanized-autonomous-organization.md:5512	kostakis bauwens 2014	Reference-list placement provides bounded context on P2P infrastructures and commons scenarios.	general reference	Publisher contents, chapters 3–4, establish P2P infrastructure and commons scenarios, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B002-O003	removed	docs/13-member-trust-union/03-governance-and-risk-distribution.md:151	bauwens kostakis 2014	The pre-audit reference reversed the title-page author order, listing Bauwens before Kostakis.	general reference	The Springer Nature record identifies Vasilis Kostakis as first author and Michel Bauwens as second author.	Restore the verified title-page author order while retaining the book as scoped context.
B002-O004	removed	docs/humanized-autonomous-organization.md:5522	bauwens kostakis 2014	The consolidated pre-audit reference reversed the title-page author order, listing Bauwens before Kostakis.	general reference	The Springer Nature record identifies Vasilis Kostakis as first author and Michel Bauwens as second author.	Restore the verified title-page author order while retaining the book as scoped context.

B003 — Commons-based peer production and the economics of the commons

- **Canonical citation:** No canonical publication was established for the baseline Bauwens and Kostakis (2015) title.
- **Metadata status:** not corroborated; no first-party record for the baseline title was located, and the verified 2015 Bauwens–Kostakis article is a distinct work
- **Primary record:** No primary record for the baseline identity was located. Comparison only: official *Journal of Peer Production* record for the distinct 2015 article; author-institution record for that article.
- **Disposition:** removed
- **Defect tags:** bauwens-kostakis-commons-economics
- **Forbidden live strings:** Commons-based peer production and the economics of the commons

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B003-O001	removed	docs/04-economics/02-revenue-allocation-framework.md:95	bauwens kostakis 2015	Pre-audit reference to the uncorroborated Bauwens and Kostakis (2015) title.	not verifiable	The search could not locate a primary record in the journal, author archive, P2P Foundation, Commons Strategies Group, or DOI channels. The distinct 2015 article does not establish this identity or the chapter's allocation formulas.	Remove this entry. Do not transfer its placement to the distinct article; audit any future addition as a newly named, scoped reference.
B003-O002	removed	docs/humanized-autonomous-organization.md:834	bauwens kostakis 2015	Consolidated duplicate of the uncorroborated Bauwens and Kostakis (2015) title.	not verifiable	The search could not locate a primary record in the journal, author archive, P2P Foundation, Commons Strategies Group, or DOI channels. The distinct 2015 article does not establish this identity or the chapter's allocation formulas.	Remove this entry. Do not transfer its placement to the distinct article; audit any future addition as a newly named, scoped reference.

B004 — Peer to Peer: The Commons Manifesto

- **Canonical citation:** Bauwens, M., Kostakis, V., & Pazaitis, A. (2019). *Peer to Peer: The Commons Manifesto*. University of Westminster Press. <https://doi.org/10.16997/book33>
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** University of Westminster Press record
- **Disposition:** retained
- **Defect tags:** commons-manifesto-coauthor-publisher
- **Forbidden live strings:** Bauwens, M., & Pazaitis, A. (2019)

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B004-O001	live	docs/13-member-trust-union/04-technology-and-physical-digital-integration.md:139	bauwens kostakis pazaitis 2019	Reference-list placement provides bounded context on P2P social relations and technical infrastructures.	general reference	Publisher abstract and DOI establish P2P as social relation and technical infrastructure, not the corpus's stack.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B004-O002	live	docs/13-member-trust-union/06-integration-with-icn-and-external-markets.md:145	bauwens kostakis pazaitis 2019	Reference-list placement provides bounded context on P2P social relations and technical infrastructures.	general reference	Publisher abstract and DOI establish P2P as social relation and technical infrastructure, not the corpus's stack.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B004-O003	live	docs/13-member-trust-union/08-mtu-as-domain-specific-hao-implementation.md:121	bauwens kostakis pazaitis 2019	Reference-list placement provides bounded context on P2P social relations and technical infrastructures.	general reference	Publisher abstract and DOI establish P2P as social relation and technical infrastructure, not the corpus's stack.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B004-O004	live	docs/humanized-autonomous-organization.md:5659	bauwens kostakis pazaitis 2019	Reference-list placement provides bounded context on P2P social relations and technical infrastructures.	general reference	Publisher abstract and DOI establish P2P as social relation and technical infrastructure, not the corpus's stack.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B004-O005	live	docs/humanized-autonomous-organization.md:5938	bauwens kostakis pazaitis 2019	Reference-list placement provides bounded context on P2P social relations and technical infrastructures.	general reference	Publisher abstract and DOI establish P2P as social relation and technical infrastructure, not the corpus's stack.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B004-O006	live	docs/humanized-autonomous-organization.md:6208	bauwens kostakis pazaitis 2019	Reference-list placement provides bounded context on P2P social relations and technical infrastructures.	general reference	Publisher abstract and DOI establish P2P as social relation and technical infrastructure, not the corpus's stack.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B004-O007	removed	docs/13-member-trust-union/04-technology-and-physical-digital-integration.md:139	bauwens pazaitis 2019	The pre-audit reference omitted coauthor Vasilis Kostakis and named Westminster University Press rather than University of Westminster Press.	general reference	The publisher record and DOI identify Bauwens, Kostakis, and Pazaitis as coauthors and University of Westminster Press as publisher.	Restore the full author list and canonical publisher while retaining the book as scoped context.
B004-O008	removed	docs/13-member-trust-union/06-integration-with-icn-and-external-markets.md:145	bauwens pazaitis 2019	The pre-audit reference omitted coauthor Vasilis Kostakis and named Westminster University Press rather than University of Westminster Press.	general reference	The publisher record and DOI identify Bauwens, Kostakis, and Pazaitis as coauthors and University of Westminster Press as publisher.	Restore the full author list and canonical publisher while retaining the book as scoped context.
B004-O009	removed	docs/13-member-trust-union/08-mtu-as-domain-specific-hao-implementation.md:123	bauwens pazaitis 2019	The pre-audit reference omitted coauthor Vasilis Kostakis and named Westminster University Press rather than University of Westminster Press.	general reference	The publisher record and DOI identify Bauwens, Kostakis, and Pazaitis as coauthors and University of Westminster Press as publisher.	Restore the full author list and canonical publisher while retaining the book as scoped context.
B004-O010	removed	docs/humanized-autonomous-organization.md:5671	bauwens pazaitis 2019	The consolidated pre-audit reference omitted coauthor Vasilis Kostakis and named Westminster University Press rather than University of Westminster Press.	general reference	The publisher record and DOI identify Bauwens, Kostakis, and Pazaitis as coauthors and University of Westminster Press as publisher.	Restore the full author list and canonical publisher while retaining the book as scoped context.
B004-O011	removed	docs/humanized-autonomous-organization.md:5950	bauwens pazaitis 2019	The consolidated pre-audit reference omitted coauthor Vasilis Kostakis and named Westminster University Press rather than University of Westminster Press.	general reference	The publisher record and DOI identify Bauwens, Kostakis, and Pazaitis as coauthors and University of Westminster Press as publisher.	Restore the full author list and canonical publisher while retaining the book as scoped context.
B004-O012	removed	docs/humanized-autonomous-organization.md:6225	bauwens pazaitis 2019	The consolidated pre-audit reference omitted coauthor Vasilis Kostakis and named Westminster University Press rather than University of Westminster Press.	general reference	The publisher record and DOI identify Bauwens, Kostakis, and Pazaitis as coauthors and University of Westminster Press as publisher.	Restore the full author list and canonical publisher while retaining the book as scoped context.

B005 — The Commons Transition Plan

- **Canonical citation:** No canonical 2021 single-author publication was established for the baseline string.
- **Metadata status:** not corroborated; the official catalogue instead identifies distinct 2014 Ecuador and 2017 Ghent commons-transition plans
- **Primary record:** No primary record for the baseline 2021 single-author identity was located. Comparison only: [official P2P Foundation publication catalogue](#), which separately lists *The FLOK Report: A Commons Transition Plan for the State of Ecuador* and *Commons Transition Plan for the City of Ghent*.
- **Disposition:** removed
- **Defect tags:** bauwens-commons-transition
- **Forbidden live strings:** The Commons Transition Plan

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B005-O001	removed	docs/04-economics/05-reinvestment-mechanisms.md:149	bauwens 2021	Pre-audit reference to an uncorroborated 2021 single-author commons-transition plan.	not verifiable	The search could not locate the baseline identity; the official catalogue records differently authored and dated Ecuador and Ghent plans, not the stated work or the chapter's reinvestment protocols.	Remove the entry. Do not substitute either real plan on title similarity; audit any future citation under its actual identity.
B005-O002	removed	docs/humanized-autonomous-organization.md:1202	bauwens 2021	Consolidated duplicate of the uncorroborated 2021 single-author commons-transition plan.	not verifiable	The search could not locate the baseline identity; the official catalogue records differently authored and dated Ecuador and Ghent plans, not the stated work or the chapter's reinvestment protocols.	Remove the entry. Do not substitute either real plan on title similarity; audit any future citation under its actual identity.

B006 — Designing Investor-Compatible Commons-Based Models

- **Canonical citation:** No canonical publication was established for the baseline string.
- **Metadata status:** not corroborated; Bauwens and Niaros (2017) is a distinct candidate work, not the baseline publication
- **Primary record:** No primary record for the baseline title was located. Comparison only: [official P2P Foundation catalogue entry](#) and [Heinrich Böll Foundation co-publisher record](#) for the distinct Bauwens and Niaros (2017) report.
- **Disposition:** removed
- **Defect tags:** bauwens-investor-compatible
- **Forbidden live strings:** Designing Investor-Compatible Commons-Based Models

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B006-O001	removed	docs/04-economics/06-external-interface-economics.md:131	bauwens 2021	Pre-audit reference to an uncorroborated investor-compatible commons title.	not verifiable	The search could not locate the baseline record. The distinct 2017 report, pp. 23–25, discusses internal value regimes and reciprocity licensing, not the chapter's caps, schedules, or market firewalls.	Remove the invalid entry. Do not introduce the 2017 report in this repair; audit it separately if later added.
B006-O002	removed	docs/humanized-autonomous-organization.md:1339	bauwens 2021	Consolidated duplicate of the uncorroborated investor-compatible commons title.	not verifiable	The search could not locate the baseline record. The distinct 2017 report, pp. 23–25, discusses internal value regimes and reciprocity licensing, not the chapter's caps, schedules, or market firewalls.	Remove the invalid entry. Do not introduce the 2017 report in this repair; audit it separately if later added.

B007 — The Partner State & the Commons Economy

- **Canonical citation:** No canonical 2021 publication was established for the baseline string.
- **Metadata status:** not corroborated; the accessible partner-state texts are distinct works and do not establish the baseline 2021 identity
- **Primary record:** No primary record for the baseline identity was located. Comparison only: official Commons Transition Primer text for *Blueprint for a Partner State* and official journal record for the distinct 2015 Bauwens–Kostakis article.
- **Disposition:** removed
- **Defect tags:** bauwens-partner-state
- **Forbidden live strings:** The Partner State & the Commons Economy

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B007-O001	removed	docs/04-economics/03-sep-economic-agreements.md:93	bauwens 2021	Pre-audit reference to an uncorroborated 2021 partner-state title.	not verifiable	The search could not locate the stated identity. Related primary texts support a general partner-state frame but not this record or the chapter's SEP allocations, lifecycle, and exit rules.	Remove the entry. Do not silently substitute a related essay; audit any future candidate under its actual metadata.
B007-O002	removed	docs/humanized-autonomous-organization.md:927	bauwens 2021	Consolidated duplicate of the uncorroborated 2021 partner-state title.	not verifiable	The search could not locate the stated identity. Related primary texts support a general partner-state frame but not this record or the chapter's SEP allocations, lifecycle, and exit rules.	Remove the entry. Do not silently substitute a related essay; audit any future candidate under its actual metadata.

B008 — Restoring Sanctuary: A New Operating System for Trauma-Informed Systems of Care

- **Canonical citation:** Bloom, S. L., & Farragher, B. (2013). *Restoring Sanctuary: A New Operating System for Trauma-Informed Systems of Care*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199796366.001.0001>
- **Metadata status:** verified against the linked primary record
- **Primary record:** Oxford Academic book record
- **Disposition:** retained
- **Defect tags:** bloom-farragher-inline-only
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B008-O001	live	docs/07-human-systems/00-overview.md:27	bloom farragher 2013	Reference-list placement provides bounded context on trauma-informed organizational change.	general reference	Oxford chapter 2 and chapter 4 abstracts describe trauma-informed organizational change and healthy democratic functioning.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B008-O002	live	docs/07-human-systems/00-overview.md:9	bloom farragher 2013	Trauma-informed organizational change is one input to the human-systems synthesis.	supported	Oxford chapter 2 and chapter 4 abstracts describe trauma-informed organizational change and healthy democratic functioning.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B008-O003	live	docs/humanized-autonomous-organization.md:1977	bloom farragher 2013	Trauma-informed organizational change is one input to the human-systems synthesis.	supported	Oxford chapter 2 and chapter 4 abstracts describe trauma-informed organizational change and healthy democratic functioning.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B008-O004	live	docs/humanized-autonomous-organization.md:1995	bloom farragher 2013	Reference-list placement provides bounded context on trauma-informed organizational change.	general reference	Oxford chapter 2 and chapter 4 abstracts describe trauma-informed organizational change and healthy democratic functioning.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B008-O005	removed	docs/07-human-systems/00-overview.md:9	bloom farragher 2013	The pre-audit sentence treated trauma-informed systems design as support for HAO resilience and inclusion as the organization grows.	partially supported	Oxford chapter 2 and chapter 4 abstracts establish trauma-informed organizational change and healthy democratic functioning, not outcomes for HAOs or organizational scaling.	Narrow the source-backed influence and identify the HAO outcomes as design intentions.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B008-O006	removed	docs/humanized-autonomous-organization.md:1983	bloom farragher 2013	The consolidated pre-audit sentence said the cited disciplines ensure HAO resilience, inclusion, and meaning generation as the organization grows.	partially supported	Oxford chapter 2 and chapter 4 abstracts establish trauma-informed organizational change and healthy democratic functioning, not guaranteed HAO outcomes or scaling effects.	Replace the guarantee with a bounded influence statement and corpus-owned design intention.

B009 — Commons Engine design-pattern titles

- **Canonical citation:** No canonical publication was established for either Commons Engine (2020) title in the baseline.
- **Metadata status:** not currently verifiable; neither alleged 2020 title was located in the bounded first-party search, and intermittent errors limited access to some Commons Engine routes
- **Primary record:** No primary record for either alleged publication was located. Commons Engine's first-party site and its about, services, and blog/archive routes were checked as the bounded search surface; the homepage is an organizational comparison source, not a publication record.
- **Disposition:** removed
- **Defect tags:** commons-engine-title
- **Forbidden live strings:** Design Patterns for Cooperatives in the Digital Economy; Design Patterns for Commons-Oriented Governance

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B009-O001	removed	docs/13-member-trust-union/06-integration-with-icn-and-external-markets.md:151	commons engine 2020	Pre-audit reference to <i>Design Patterns for Cooperatives in the Digital Economy</i> .	not verifiable	The search could not locate an official record or full text. Reachable first-party material did not establish this publication or the surrounding market mechanisms; some Commons Engine routes intermittently returned access errors.	Remove the citation, not replace it with the homepage. Restore only if an archived first-party publication is obtained and audited.
B009-O002	removed	docs/13-member-trust-union/07-ethical-and-regulatory-anchoring.md:144	commons engine 2020	Pre-audit reference to <i>Design Patterns for Commons-Oriented Governance</i> .	not verifiable	The search could not locate an official record or full text. Reachable first-party material did not establish this publication or the surrounding regulatory mechanisms; some Commons Engine routes intermittently returned access errors.	Remove the citation, not replace it with the homepage. Restore only if an archived first-party publication is obtained and audited.
B009-O003	removed	docs/13-member-trust-union/08-mtu-as-domain-specific-hao-implementation.md:119	commons engine 2020	Pre-audit reference to <i>Design Patterns for Cooperatives in the Digital Economy</i> .	not verifiable	The search could not locate an official record or full text. Reachable first-party material did not establish this publication or the surrounding implementation mechanisms; some Commons Engine routes intermittently returned access errors.	Remove the citation, not replace it with the homepage. Restore only if an archived first-party publication is obtained and audited.
B009-O004	removed	docs/humanized-autonomous-organization.md:5956	commons engine 2020	Consolidated pre-audit reference to <i>Design Patterns for Cooperatives in the Digital Economy</i> .	not verifiable	The search could not locate an official record or full text. Reachable first-party material did not establish this publication or the surrounding market mechanisms; some Commons Engine routes intermittently returned access errors.	Remove the citation, not replace it with the homepage. Restore only if an archived first-party publication is obtained and audited.
B009-O005	removed	docs/humanized-autonomous-organization.md:6100	commons engine 2020	Consolidated pre-audit reference to <i>Design Patterns for Commons-Oriented Governance</i> .	not verifiable	The search could not locate an official record or full text. Reachable first-party material did not establish this publication or the surrounding regulatory mechanisms; some Commons Engine routes intermittently returned access errors.	Remove the citation, not replace it with the homepage. Restore only if an archived first-party publication is obtained and audited.
B009-O006	removed	docs/humanized-autonomous-organization.md:6221	commons engine 2020	Consolidated pre-audit reference to <i>Design Patterns for Cooperatives in the Digital Economy</i> .	not verifiable	The search could not locate an official record or full text. Reachable first-party material did not establish this publication or the surrounding implementation mechanisms; some Commons Engine routes intermittently returned access errors.	Remove the citation, not replace it with the homepage. Restore only if an archived first-party publication is obtained and audited.

B010 — Democratic Money and Capital for the Commons

- **Canonical citation:** Bollier, D., & Conaty, P. (2016). *Democratic Money and Capital for the Commons: Strategies for Transforming Neoliberal Finance Through Commons-Based Alternatives*. Commons Strategies Group in cooperation with the Heinrich Böll Foundation.
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** Heinrich Böll Foundation publication record
- **Disposition:** retained
- **Defect tags:** democratic-money-authorship-year
- **Forbidden live strings:** Commons Strategies Group (2015). *Democratic Money and Capital for the Commons*

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B010-O001	live	docs/04-economics/06-external-interface-economics.md:128	bollier conaty 2016	Reference-list placement provides bounded context on commons-based finance.	general reference	Official report §II-III concerns commons-based finance and institutional forms, not HAO investor interfaces.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B010-0002	live	docs/humanized-autonomous-organization.md:1331	bollier conaty 2016	Reference-list placement provides bounded context on commons-based finance.	general reference	Official report §II-III concerns commons-based finance and institutional forms, not HAO investor interfaces.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B010-0003	removed	docs/04-economics/06-external-interface-economics.md:128	commons strategies group 2015	The pre-audit reference treated Commons Strategies Group as author and dated the report 2015.	general reference	The Heinrich Böll Foundation record identifies David Bollier and Pat Conaty as authors and publishes the report in January 2016.	Replace the workshop-era group/year attribution with the verified authors and publication year.
B010-0004	removed	docs/humanized-autonomous-organization.md:1336	commons strategies group 2015	The consolidated pre-audit reference treated Commons Strategies Group as author and dated the report 2015.	general reference	The Heinrich Böll Foundation record identifies David Bollier and Pat Conaty as authors and publishes the report in January 2016.	Replace the workshop-era group/year attribution with the verified authors and publication year.

B011 — Patterns of Commoning

- **Canonical citation:** Bollier, D., & Helfrich, S. (Eds.). (2015). *Patterns of Commoning*. Commons Strategies Group in cooperation with Off the Common Books.
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** [Commons Strategies Group publication page](#)
- **Disposition:** retained
- **Defect tags:** patterns-commoning-editor-role
- **Forbidden live strings:** Commons Strategies Group. (2015). *Patterns of Commoning*

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B011-0001	live	docs/04-economics/03-sep-economic-agreements.md:96	bollier helfrich 2015	Reference-list placement provides bounded context on commoning practices and collective governance.	general reference	Official contents and Overture section establish commoning and collective governance, not SEP formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B011-0002	live	docs/humanized-autonomous-organization.md:929	bollier helfrich 2015	Reference-list placement provides bounded context on commoning practices and collective governance.	general reference	Official contents and Overture section establish commoning and collective governance, not SEP formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B011-0003	removed	docs/04-economics/03-sep-economic-agreements.md:97	commons strategies group 2015	The pre-audit reference treated Commons Strategies Group as author rather than naming Bollier and Helfrich as editors.	general reference	The publication page and volume front matter identify David Bollier and Silke Helfrich as editors; Commons Strategies Group is the publishing organization.	Restore editor responsibility while retaining the volume as scoped context.
B011-0004	removed	docs/humanized-autonomous-organization.md:931	commons strategies group 2015	The consolidated pre-audit reference treated Commons Strategies Group as author rather than naming Bollier and Helfrich as editors.	general reference	The publication page and volume front matter identify David Bollier and Silke Helfrich as editors; Commons Strategies Group is the publishing organization.	Restore editor responsibility while retaining the volume as scoped context.

B012 — Common: On Revolution in the 21st Century

- **Canonical citation:** Dardot, P., & Laval, C. (2019). *Common: On Revolution in the 21st Century* (M. MacLellan, Trans.; I. Szeman, Preface). Bloomsbury Academic.
- **Metadata status:** verified against the linked primary record
- **Primary record:** [Bloomsbury Academic book record](#)
- **Disposition:** retained
- **Defect tags:** none
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B012-0001	live	docs/13-member-trust-union/05-member-experience-and-compensation.md:134	dardot laval 2019	Reference-list placement provides bounded context on the common as collective practice and institution.	general reference	Publisher abstract and contents establish collective practice and shared rules, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B012-0002	live	docs/13-member-trust-union/07-ethical-and-regulatory-anchoring.md:140	dardot laval 2019	Reference-list placement provides bounded context on the common as collective practice and institution.	general reference	Publisher abstract and contents establish collective practice and shared rules, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B012-0003	live	docs/13-member-trust-union/08-mtu-as-domain-specific-hao-implementation.md:119	dardot laval 2019	Reference-list placement provides bounded context on the common as collective practice and institution.	general reference	Publisher abstract and contents establish collective practice and shared rules, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B012-0004	live	docs/humanized-autonomous-organization.md:5793	dardot laval 2019	Reference-list placement provides bounded context on the common as collective practice and institution.	general reference	Publisher abstract and contents establish collective practice and shared rules, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

Occurrence State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B012-O005 live	docs/humanized-autonomous-organization.md:6082	dardot laval 2019	Reference-list placement provides bounded context on the common as collective practice and institution.	general reference	Publisher abstract and contents establish collective practice and shared rules, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B012-O006 live	docs/humanized-autonomous-organization.md:6206	dardot laval 2019	Reference-list placement provides bounded context on the common as collective practice and institution.	general reference	Publisher abstract and contents establish collective practice and shared rules, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

B013 — Blockchain and the Law: The Rule of Code

- **Canonical citation:** De Filippi, P., & Wright, A. (2018). *Blockchain and the Law: The Rule of Code*. Harvard University Press. <https://doi.org/10.4159/9780674985933>
- **Metadata status:** verified against the linked primary record
- **Primary record:** [Harvard University Press DOI record](#)
- **Disposition:** retained
- **Defect tags:** none
- **Forbidden live strings:** none

Occurrence State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B013-O001 live	docs/13-member-trust-union/07-ethical-and-regulatory-anchoring.md:142	de filippi wright 2018	Reference-list placement provides bounded context on blockchain architecture and law.	general reference	Publisher contents, pp. 173 and 193, establish a relationship between code, blockchain, and regulation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B013-O002 live	docs/humanized-autonomous-organization.md:6084	de filippi wright 2018	Reference-list placement provides bounded context on blockchain architecture and law.	general reference	Publisher contents, pp. 173 and 193, establish a relationship between code, blockchain, and regulation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

B014 — Systems Thinking and Sustainability

- **Canonical citation:** No canonical Donella Meadows Institute (2015) publication was established for the baseline string.
- **Metadata status:** not corroborated; the official archive contains a related 2015 post, but not the baseline title
- **Primary record:** No primary record for the baseline identity was located. Comparison only: [official Donella Meadows Project resource archive](#) and the distinct 2015 post *Coming Back to Our Systems Roots*.
- **Disposition:** removed
- **Defect tags:** meadows-systems-thinking
- **Forbidden live strings:** Systems Thinking and Sustainability

Occurrence State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B014-O001 removed	docs/04-economics/05-reinvestment-mechanisms.md:152	donella meadows institute 2015	Pre-audit reference to the uncorroborated <i>Systems Thinking and Sustainability</i> title.	not verifiable	The official archive search could not locate the stated identity or full text. Its distinct 2015 post offers general systems commentary but not the chapter's reinvestment pools or audit mechanisms.	Remove the baseline reference. Do not substitute the related post without a separately audited, narrower claim.
B014-O002 removed	docs/humanized-autonomous-organization.md:1205	donella meadows institute 2015	Consolidated duplicate of the uncorroborated <i>Systems Thinking and Sustainability</i> title.	not verifiable	The official archive search could not locate the stated identity or full text. Its distinct 2015 post offers general systems commentary but not the chapter's reinvestment pools or audit mechanisms.	Remove the baseline reference. Do not substitute the related post without a separately audited, narrower claim.

B015 — Cannibals with Forks

- **Canonical citation:** Elkington, J. (1999). *Cannibals with Forks: The Triple Bottom Line of 21st Century Business*. Capstone Publishing.
- **Metadata status:** verified; current corpus placement removed after claim review
- **Primary record:** [Authoritative 1999 paperback catalog record](#); [Wiley publisher abstract used for claim review](#)
- **Disposition:** removed
- **Defect tags:** none
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B015-0001	removed	docs/13-member-trust-union/03-governance-and-risk-distribution.md:155	elkington 1999	Historical pre-audit occurrence using the superseded or removed elkington 1999 attribution.	unsupported	The verified primary source's abstract establishes a different subject and no identifiable support for this placement.	Remove the unsupported placement while preserving the work's real identity in this ledger.
B015-0002	removed	docs/humanized-autonomous-organization.md:5526	elkington 1999	Historical pre-audit occurrence using the superseded or removed elkington 1999 attribution.	unsupported	The verified primary source's abstract establishes a different subject and no identifiable support for this placement.	Remove the unsupported placement while preserving the work's real identity in this ledger.

B016 — Three Strategic Concepts for the Guidance of Co-operatives

- **Canonical citation:** Fairbairn, B. (2003). *Three Strategic Concepts for the Guidance of Co-operatives: Linkage, Transparency, and Cognition*. Centre for the Study of Co-operatives, University of Saskatchewan. <https://doi.org/10.22004/ag.econ.31755>
- **Metadata status:** verified; pre-audit title completed
- **Primary record:** [AgEcon Search institutional record](#)
- **Disposition:** retained
- **Defect tags:** fairbairn-title
- **Forbidden live strings:** Fairbairn, B. (2003). *Three Strategic Concepts for the Guidance of Co-operatives*

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B016-0001	live	docs/04-economics/05-reinvestment-mechanisms.md:151	fairbairn 2003	Reference-list placement provides bounded context on cooperative linkage, surplus, and internal capitalization.	general reference	Original report p. 27 discusses internal capitalization from cooperative surplus and investment cognition.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B016-0002	live	docs/humanized-autonomous-organization.md:1201	fairbairn 2003	Reference-list placement provides bounded context on cooperative linkage, surplus, and internal capitalization.	general reference	Original report p. 27 discusses internal capitalization from cooperative surplus and investment cognition.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B016-0003	removed	docs/04-economics/05-reinvestment-mechanisms.md:153	fairbairn 2003	The pre-audit reference truncated the title after <i>Three Strategic Concepts for the Guidance of Co-operatives</i> .	general reference	The institutional repository record and original report identify the subtitle <i>Linkage, Transparency, and Cognition</i> and the University of Saskatchewan series.	Complete the title and publisher information while retaining the work as scoped context.
B016-0004	removed	docs/humanized-autonomous-organization.md:1206	fairbairn 2003	Consolidated pre-audit duplicate used the truncated title and omitted publisher information.	general reference	The institutional repository record and original report identify the subtitle <i>Linkage, Transparency, and Cognition</i> and the University of Saskatchewan series.	Complete the title and publisher information while retaining the work as scoped context.

B017 — Slicing Pie: Funding Your Business Without Funds

- **Canonical citation:** Moyer, M. (2012). *Slicing Pie: Funding Your Business Without Funds*. Lake Shark Ventures.
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** [Official author-hosted 2012 sample](#)
- **Disposition:** retained
- **Defect tags:** slicing-pie-author
- **Forbidden live strings:** Fairfield, M. (2012) ; (Fairfield, 2012) ; Funding Your Company Without Funds

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B017-0001	live	docs/04-economics/02-revenue-allocation-framework.md:39	moyer 2012	Risk-adjusted, contribution-proportional slices inform the corpus's optional dynamic-equity design.	supported	Official sample and model pages, pp. 1-15 and model sections, tie ownership to risk-adjusted time, cash, and resources.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B017-0002	live	docs/04-economics/02-revenue-allocation-framework.md:92	moyer 2012	Reference-list placement provides bounded context on risk-adjusted dynamic equity.	general reference	Official sample and model pages, pp. 1-15 and model sections, tie ownership to risk-adjusted time, cash, and resources.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B017-0003	live	docs/04-economics/04-equity-and-member-compensation.md:117	moyer 2012	Reference-list placement provides bounded context on risk-adjusted dynamic equity.	general reference	Official sample and model pages, pp. 1-15 and model sections, tie ownership to risk-adjusted time, cash, and resources.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B017-0004	live	docs/04-economics/04-equity-and-member-compensation.md:18	moyer 2012	Risk-adjusted, contribution-proportional slices inform the corpus's optional dynamic-equity design.	supported	Official sample and model pages, pp. 1-15 and model sections, tie ownership to risk-adjusted time, cash, and resources.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B017-0005	live	docs/04-economics/04-equity-and-member-compensation.md:41	moyer 2012	Risk-adjusted, contribution-proportional slices inform the corpus's optional dynamic-equity design.	supported	Official sample and model pages, pp. 1-15 and model sections, tie ownership to risk-adjusted time, cash, and resources.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B017-O006	live	docs/12-icn/03-economic-model.md:99	moyer 2012	Risk-adjusted, contribution-proportional slices inform the corpus's optional dynamic-equity design.	supported	Official sample and model pages, pp. 1-15 and model sections, tie ownership to risk-adjusted time, cash, and resources.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B017-O007	live	docs/humanized-autonomous-organization.md:1046	moyer 2012	Reference-list placement provides bounded context on risk-adjusted dynamic equity.	general reference	Official sample and model pages, pp. 1-15 and model sections, tie ownership to risk-adjusted time, cash, and resources.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B017-O008	live	docs/humanized-autonomous-organization.md:4619	moyer 2012	Risk-adjusted, contribution-proportional slices inform the corpus's optional dynamic-equity design.	supported	Official sample and model pages, pp. 1-15 and model sections, tie ownership to risk-adjusted time, cash, and resources.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B017-O009	live	docs/humanized-autonomous-organization.md:776	moyer 2012	Risk-adjusted, contribution-proportional slices inform the corpus's optional dynamic-equity design.	supported	Official sample and model pages, pp. 1-15 and model sections, tie ownership to risk-adjusted time, cash, and resources.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B017-O010	live	docs/humanized-autonomous-organization.md:831	moyer 2012	Reference-list placement provides bounded context on risk-adjusted dynamic equity.	general reference	Official sample and model pages, pp. 1-15 and model sections, tie ownership to risk-adjusted time, cash, and resources.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B017-O011	live	docs/humanized-autonomous-organization.md:947	moyer 2012	Risk-adjusted, contribution-proportional slices inform the corpus's optional dynamic-equity design.	supported	Official sample and model pages, pp. 1-15 and model sections, tie ownership to risk-adjusted time, cash, and resources.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B017-O012	live	docs/humanized-autonomous-organization.md:970	moyer 2012	Risk-adjusted, contribution-proportional slices inform the corpus's optional dynamic-equity design.	supported	Official sample and model pages, pp. 1-15 and model sections, tie ownership to risk-adjusted time, cash, and resources.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B017-O013	removed	docs/04-economics/02-revenue-allocation-framework.md:39	fairfield 2012	A Slicing Pie-style model allocates equity from risk-adjusted time and resource contributions.	supported	The official sample and model sections, pp. 1-15, tie ownership to each participant's share of total at-risk contributions.	Correct the author to Moyer and retain the bounded dynamic-equity claim.
B017-O014	removed	docs/04-economics/02-revenue-allocation-framework.md:92	fairfield 2012	Pre-audit reference-list placement for the title-identifiable Slicing Pie work.	general reference	The official 2012 title and copyright pages identify Moyer's book and its risk-adjusted dynamic-equity subject.	Correct the author and title while retaining the work as scoped context.
B017-O015	removed	docs/04-economics/04-equity-and-member-compensation.md:117	fairfield 2012	Pre-audit reference-list placement for the title-identifiable Slicing Pie work.	general reference	The official 2012 title and copyright pages identify Moyer's book and its risk-adjusted dynamic-equity subject.	Correct the author and title while retaining the work as scoped context.
B017-O016	removed	docs/04-economics/04-equity-and-member-compensation.md:18	fairfield 2012	The pre-audit sentence treated dynamic equity as support for the chapter's full four-part principles.	partially supported	The official sample and model sections, pp. 1-15, support dynamic contribution-proportional equity, not the corpus's complete multi-form, portable, and temporal framework.	Correct the author and limit the attribution to dynamic contribution-proportional equity.
B017-O017	removed	docs/04-economics/04-equity-and-member-compensation.md:41	fairfield 2012	The pre-audit block said most UMEs and SEPs use a Slicing Pie model and bundled source-backed inputs with HAO-specific recovery tiers.	partially supported	Official model sections, pp. 1-15, support at-risk inputs, rolling proportions, multipliers, and departure categories, but not HAO adoption or trust-gated recovery.	Correct the author, make adoption optional, and label trust-gated recovery as a corpus extension.
B017-O018	removed	docs/humanized-autonomous-organization.md:1048	fairfield 2012	Consolidated pre-audit reference-list placement for the title-identifiable Slicing Pie work.	general reference	The official 2012 title and copyright pages identify Moyer's book and its risk-adjusted dynamic-equity subject.	Correct the author and title while retaining the work as scoped context.
B017-O019	removed	docs/humanized-autonomous-organization.md:776	fairfield 2012	Consolidated pre-audit claim that a Slicing Pie-style model allocates equity from risk-adjusted contributions.	supported	The official sample and model sections, pp. 1-15, tie ownership to each participant's share of total at-risk contributions.	Correct the author to Moyer and retain the bounded dynamic-equity claim.
B017-O020	removed	docs/humanized-autonomous-organization.md:831	fairfield 2012	Consolidated pre-audit reference-list placement for the title-identifiable Slicing Pie work.	general reference	The official 2012 title and copyright pages identify Moyer's book and its risk-adjusted dynamic-equity subject.	Correct the author and title while retaining the work as scoped context.
B017-O021	removed	docs/humanized-autonomous-organization.md:949	fairfield 2012	Consolidated pre-audit attribution of the chapter's full four-part principles to dynamic equity theory.	partially supported	The official sample and model sections, pp. 1-15, support dynamic contribution-proportional equity, not the corpus's complete multi-form, portable, and temporal framework.	Correct the author and limit the attribution to dynamic contribution-proportional equity.
B017-O022	removed	docs/humanized-autonomous-organization.md:972	fairfield 2012	Consolidated pre-audit block asserting HAO adoption and bundling source-backed model elements with trust-gated recovery.	partially supported	Official model sections, pp. 1-15, support at-risk inputs, rolling proportions, multipliers, and departure categories, but not HAO adoption or trust-gated recovery.	Correct the author, make adoption optional, and label trust-gated recovery as a corpus extension.
B017-O023	removed	docs/12-icn/03-economic-model.md:99	slicing pie	The uncited pre-audit ICN claim said equity is calculated using Slicing Pie slices for time, money, IP, or network value.	partially supported	Official model sections, pp. 1-15, support varied at-risk contributions and dynamic slices, but not ICN adoption or automatic treatment of network value.	Cite Moyer, make adoption optional, and identify network value as a corpus extension.
B017-O024	removed	docs/humanized-autonomous-organization.md:4619	slicing pie	Consolidated duplicate of the uncited pre-audit ICN adoption and network-value claim.	partially supported	Official model sections, pp. 1-15, support varied at-risk contributions and dynamic slices, but not ICN adoption or automatic treatment of network value.	Cite Moyer, make adoption optional, and identify network value as a corpus extension.

B018 — The Evidence-Based Story of Savings Groups

- **Canonical citation:** Gash, M., & Odell, K. (2013). *The Evidence-Based Story of Savings Groups: A Synthesis of Seven Randomized Control Trials*. SEEP Research, Savings-led Financial Services Working Group at the SEEP Network.
- **Metadata status:** verified; pre-audit unstable home-page link replaced
- **Primary record:** [Original SEEP report](#)
- **Disposition:** retained
- **Defect tags:** savings-groups-source-link
- **Forbidden live strings:** Gash, M., & Odell, K. (2013). *The Evidence-Based Story of Savings Groups: A Synthesis of Seven Randomized Control Trials*. The SEEP Network. <https://seepnetwork.org>

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B018-O001	live	docs/13-member-trust-union/02-financial-product-design.md:157	gash odell 2013	Reference-list placement provides bounded context on accumulating savings groups and share-out cycles.	general reference	Original report pp. 7–8 describes a common fund, member loans, and end-of-cycle share-out.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B018-O002	live	docs/13-member-trust-union/02-financial-product-design.md:93	gash odell 2013	Savings groups provide a distinct accumulating-fund and share-out precedent.	supported	Original report pp. 7–8 describes a common fund, member loans, and end-of-cycle share-out.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B018-O003	live	docs/humanized-autonomous-organization.md:5291	gash odell 2013	Savings groups provide a distinct accumulating-fund and share-out precedent.	supported	Original report pp. 7–8 describes a common fund, member loans, and end-of-cycle share-out.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B018-O004	live	docs/humanized-autonomous-organization.md:5357	gash odell 2013	Reference-list placement provides bounded context on accumulating savings groups and share-out cycles.	general reference	Original report pp. 7–8 describes a common fund, member loans, and end-of-cycle share-out.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B018-O005	removed	docs/13-member-trust-union/02-financial-product-design.md:93	gash odell 2013	The pre-audit sentence treated the report as joint support for ROSCAs, tontines, kibbutz-like redistribution, and the MACI design.	partially supported	The original report executive summary, pp. 7–8, describes accumulating savings groups, member loans, and end-of-cycle share-out, not rotating payout or the other bundled mechanisms.	Replace the bundled attribution with the distinct savings-group precedent and identify the remaining synthesis as corpus-authored.
B018-O006	removed	docs/humanized-autonomous-organization.md:5297	gash odell 2013	Consolidated pre-audit duplicate of the bundled ROSCA, tontine, kibbutz, and MACI attribution.	partially supported	The original report executive summary, pp. 7–8, describes accumulating savings groups, member loans, and end-of-cycle share-out, not rotating payout or the other bundled mechanisms.	Replace the bundled attribution with the distinct savings-group precedent and identify the remaining synthesis as corpus-authored.
B018-O007	removed	docs/13-member-trust-union/02-financial-product-design.md:157	gash odell 2013	The pre-audit reference linked only to the SEEP home page rather than the preserved report.	general reference	The preserved original report supplies the title, authors, date, and claim evidence; the home page did not identify this record stably.	Supersede the home-page link with the stable original-report link.
B018-O008	removed	docs/humanized-autonomous-organization.md:5363	gash odell 2013	Consolidated pre-audit duplicate linked only to the SEEP home page.	general reference	The preserved original report supplies the title, authors, date, and claim evidence; the home page did not identify this record stably.	Supersede the home-page link with the stable original-report link.

B019 — Trust in Finance: Historical Perspectives

- **Canonical citation:** The baseline identity is invalid: DOI 10.1093/cje/bew037 belongs to Paul Lewis's different 2017 article.
- **Metadata status:** baseline identity contradicted by official DOI or issue metadata
- **Primary record:** [Official DOI record](#)
- **Disposition:** removed
- **Defect tags:** ghosh-trust-finance
- **Forbidden live strings:** Trust in Finance: Historical Perspectives

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B019-O001	removed	docs/13-member-trust-union/01-multi-level-trust-architecture.md:115	ghosh 2017	Historical pre-audit occurrence using the superseded or removed ghosh 2017 attribution.	unsupported	The official DOI record assigns 10.1093/cje/bew037 to Paul Lewis's different article.	Remove; do not redirect the DOI or substitute the unrelated article.
B019-O002	removed	docs/humanized-autonomous-organization.md:5196	ghosh 2017	Historical pre-audit occurrence using the superseded or removed ghosh 2017 attribution.	unsupported	The official DOI record assigns 10.1093/cje/bew037 to Paul Lewis's different article.	Remove; do not redirect the DOI or substitute the unrelated article.

B020 — The End of Capitalism (As We Knew It)

- **Canonical citation:** Gibson-Graham, J. K. (2006). *The End of Capitalism (As We Knew It): A Feminist Critique of Political Economy* (New ed.). University of Minnesota Press.

- **Metadata status:** verified; pre-audit title and edition statement completed
- **Primary record:** [University of Minnesota Press book record](#)
- **Disposition:** retained
- **Defect tags:** end-capitalism-title-edition
- **Forbidden live strings:** The End of Capitalism (As We Knew It)*. University of Minnesota Press

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B020-O001	live	docs/13-member-trust-union/05-member-experience-and-compensation.md:128	gibson graham 2006	Reference-list placement provides bounded context on non-capitalocentric economic possibilities.	general reference	Publisher abstract challenges capitalism's presumed dominance and opens space for alternative economies.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B020-O002	live	docs/humanized-autonomous-organization.md:5787	gibson graham 2006	Reference-list placement provides bounded context on non-capitalocentric economic possibilities.	general reference	Publisher abstract challenges capitalism's presumed dominance and opens space for alternative economies.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B020-O003	removed	docs/13-member-trust-union/05-member-experience-and-compensation.md:128	gibson graham 2006	The pre-audit reference omitted <i>A Feminist Critique of Political Economy</i> and the 2006 new-edition statement.	general reference	The University of Minnesota Press record supplies the complete title and identifies the 2006 new edition.	Complete the title and edition statement while retaining the work as scoped context.
B020-O004	removed	docs/humanized-autonomous-organization.md:5799	gibson graham 2006	Consolidated pre-audit duplicate omitted the subtitle and 2006 new-edition statement.	general reference	The University of Minnesota Press record supplies the complete title and identifies the 2006 new edition.	Complete the title and edition statement while retaining the work as scoped context.

B021 — Radical Technologies: The Design of Everyday Life

- **Canonical citation:** Greenfield, A. (2017). *Radical Technologies: The Design of Everyday Life*. Verso.
- **Metadata status:** verified; pre-audit publisher form repaired
- **Primary record:** [Verso Spring 2017 catalogue, p. 42](#)
- **Disposition:** retained
- **Defect tags:** radical-technologies-publisher
- **Forbidden live strings:** Radical Technologies: The Design of Everyday Life*. Verso Books

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B021-O001	live	docs/13-member-trust-union/04-technology-and-physical-digital-integration.md:129	greenfield 2017	Reference-list placement provides bounded context on the social and political consequences of networked technologies.	general reference	Verso catalogue p. 42 relates networked technologies to power, values, and social organization.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B021-O002	live	docs/humanized-autonomous-organization.md:5649	greenfield 2017	Reference-list placement provides bounded context on the social and political consequences of networked technologies.	general reference	Verso catalogue p. 42 relates networked technologies to power, values, and social organization.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B021-O003	removed	docs/13-member-trust-union/04-technology-and-physical-digital-integration.md:129	greenfield 2017	The pre-audit reference named the 2017 publisher as Verso Books.	general reference	The original-edition record in Verso's Spring 2017 catalogue, p. 42, identifies the publisher as Verso.	Normalize the publisher while retaining the work as scoped context.
B021-O004	removed	docs/humanized-autonomous-organization.md:5661	greenfield 2017	Consolidated pre-audit duplicate used the superseded publisher form Verso Books.	general reference	The original-edition record in Verso's Spring 2017 catalogue, p. 42, identifies the publisher as Verso.	Normalize the publisher while retaining the work as scoped context.

B022 — The Tragedy of the Commons

- **Canonical citation:** Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859), 1243–1248. <https://doi.org/10.1126/science.162.3859.1243>
- **Metadata status:** verified; missing DOI repaired in one chapter/consolidated pair
- **Primary record:** [Science/AAAS article record](#)
- **Disposition:** retained
- **Defect tags:** hardin-doi
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B022-O001	live	docs/13-member-trust-union/03-governance-and-risk-distribution.md:155	hardin 1968	Reference-list placement provides bounded context on Hardin's open-access thought experiment.	general reference	Original article pp. 1244–47 supplies a historically bounded open-access dilemma, not MTU controls.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B022-O002	live	docs/13-member-trust-union/07-ethical-and-regulatory-anchoring.md:138	hardin 1968	Reference-list placement provides bounded context on Hardin's open-access thought experiment.	general reference	Original article pp. 1244–47 supplies a historically bounded open-access dilemma, not MTU controls.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B022-O003	live	docs/humanized-autonomous-organization.md:5516	hardin 1968	Reference-list placement provides bounded context on Hardin's open-access thought experiment.	general reference	Original article pp. 1244–47 supplies a historically bounded open-access dilemma, not MTU controls.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B022-O004	live	docs/humanized-autonomous-organization.md:6080	hardin 1968	Reference-list placement provides bounded context on Hardin's open-access thought experiment.	general reference	Original article pp. 1244–47 supplies a historically bounded open-access dilemma, not MTU controls.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B022-O005	removed	docs/13-member-trust-union/07-ethical-and-regulatory-anchoring.md:138	hardin 1968	The pre-audit reference omitted the article DOI.	general reference	The Science/AAAS record identifies DOI 10.1126/science.162.3859.1243.	Add the DOI while retaining the work as historically bounded general context.
B022-O006	removed	docs/humanized-autonomous-organization.md:6094	hardin 1968	Consolidated pre-audit duplicate omitted the article DOI.	general reference	The Science/AAAS record identifies DOI 10.1126/science.162.3859.1243.	Add the DOI while retaining the work as historically bounded general context.

B023 — Whiplash: How to Survive Our Faster Future

- **Canonical citation:** Ito, J., & Howe, J. (2016). *Whiplash: How to Survive Our Faster Future*. Grand Central Publishing.
- **Metadata status:** verified against the linked primary record
- **Primary record:** [Hachette publisher record](#)
- **Disposition:** retained
- **Defect tags:** none
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B023-O001	live	docs/13-member-trust-union/01-multi-level-trust-architecture.md:117	ito howe 2016	Reference-list placement provides bounded context on adaptive and decentralized organizing principles.	general reference	Publisher abstract lists emergence, resilience, and systems as adaptive principles, not trust architecture.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B023-O002	live	docs/humanized-autonomous-organization.md:5198	ito howe 2016	Reference-list placement provides bounded context on adaptive and decentralized organizing principles.	general reference	Publisher abstract lists emergence, resilience, and systems as adaptive principles, not trust architecture.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

B024 — Psychological Conditions of Personal Engagement and Disengagement at Work

- **Canonical citation:** Kahn, W. A. (1990). Psychological conditions of personal engagement and disengagement at work. *Academy of Management Journal*, 33(4), 692–724. <https://doi.org/10.5465/256287>
- **Metadata status:** verified against the linked primary record
- **Primary record:** [Academy of Management article record](#)
- **Primary-text check:** A 33-page local scan (SHA-256 db2cb87fa1715ca142d3369d0d50ec771d5e281b3679eeafd5c79a427d2630fb) places Kahn's three psychological conditions on printed p. 703 (PDF p. 12) and his definition of psychological safety on printed p. 708 (PDF p. 17). The latter is not on p. 700. The study's qualitative and descriptive analyses do not establish HAO outcomes.
- **Disposition:** retained
- **Defect tags:** kahn-inline-only
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B024-O001	live	docs/07-human-systems/00-overview.md:28	kahn 1990	Reference-list placement provides bounded context on workplace engagement and psychological safety.	general reference	Primary text identifies meaningfulness, safety and availability on printed p. 703, and defines psychological safety on p. 708.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B024-O002	live	docs/07-human-systems/00-overview.md:9	kahn 1990	Workplace engagement and psychological safety are inputs to the human-systems synthesis.	supported	Primary text identifies meaningfulness, safety and availability on printed p. 703, and defines psychological safety on p. 708.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B024-O003	live	docs/humanized-autonomous-organization.md:1977	kahn 1990	Workplace engagement and psychological safety are inputs to the human-systems synthesis.	supported	Primary text identifies meaningfulness, safety and availability on printed p. 703, and defines psychological safety on p. 708.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B024-O004	live	docs/humanized-autonomous-organization.md:1996	kahn 1990	Reference-list placement provides bounded context on workplace engagement and psychological safety.	general reference	Primary text identifies meaningfulness, safety and availability on printed p. 703, and defines psychological safety on p. 708.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B024-O005	removed	docs/07-human-systems/00-overview.md:9	kahn 1990	The pre-audit sentence labeled Kahn's work participatory organizational psychology and attached it to intended HAO resilience and inclusion.	partially supported	The Academy of Management abstract establishes workplace engagement, meaningfulness, safety, and availability, not the broader disciplinary label or HAO outcomes.	Replace the imprecise label with the supported workplace concepts and identify the HAO outcomes as design intentions.
B024-O006	removed	docs/humanized-autonomous-organization.md:1983	kahn 1990	The consolidated pre-audit sentence said the cited disciplines ensure HAO resilience, inclusion, and meaning generation as the organization grows.	partially supported	The Academy of Management abstract establishes workplace engagement, meaningfulness, safety, and availability, not guaranteed HAO outcomes or scaling effects.	Replace the guarantee with a bounded influence statement and corpus-owned design intention.

B025 — Owning Our Future

- **Canonical citation:** Kelly, M. (2012). *Owning Our Future: The Emerging Ownership Revolution: Journeys to a Generative Economy*. Berrett-Koehler Publishers.
- **Metadata status:** verified; pre-audit title variants completed
- **Primary record:** [Berrett-Koehler publisher record](#)
- **Disposition:** retained
- **Defect tags:** owning-future-title
- **Forbidden live strings:** Kelly, M. (2012). *Owning Our Future* ; Kelly, M. (2012). *Owning Our Future: The Emerging Ownership Revolution*

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B025-O001	live	docs/04-economics/03-sep-economic-agreements.md:95	kelly 2012	Reference-list placement provides bounded context on generative ownership design.	general reference	Publisher abstract describes five generative-ownership patterns; the corpus's four principles remain its synthesis.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B025-O002	live	docs/04-economics/04-equity-and-member-compensation.md:119	kelly 2012	Reference-list placement provides bounded context on generative ownership design.	general reference	Publisher abstract describes five generative-ownership patterns; the corpus's four principles remain its synthesis.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B025-O003	live	docs/04-economics/04-equity-and-member-compensation.md:18	kelly 2012	Generative ownership provides a broader tradition for the corpus's synthesized equity principles.	supported	Publisher abstract describes five generative-ownership patterns; the corpus's four principles remain its synthesis.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B025-O004	live	docs/04-economics/05-reinvestment-mechanisms.md:150	kelly 2012	Reference-list placement provides bounded context on generative ownership design.	general reference	Publisher abstract describes five generative-ownership patterns; the corpus's four principles remain its synthesis.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B025-O005	live	docs/04-economics/06-external-interface-economics.md:126	kelly 2012	Reference-list placement provides bounded context on generative ownership design.	general reference	Publisher abstract describes five generative-ownership patterns; the corpus's four principles remain its synthesis.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B025-O006	live	docs/humanized-autonomous-organization.md:1048	kelly 2012	Reference-list placement provides bounded context on generative ownership design.	general reference	Publisher abstract describes five generative-ownership patterns; the corpus's four principles remain its synthesis.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B025-O007	live	docs/humanized-autonomous-organization.md:1200	kelly 2012	Reference-list placement provides bounded context on generative ownership design.	general reference	Publisher abstract describes five generative-ownership patterns; the corpus's four principles remain its synthesis.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B025-O008	live	docs/humanized-autonomous-organization.md:1329	kelly 2012	Reference-list placement provides bounded context on generative ownership design.	general reference	Publisher abstract describes five generative-ownership patterns; the corpus's four principles remain its synthesis.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B025-O009	live	docs/humanized-autonomous-organization.md:928	kelly 2012	Reference-list placement provides bounded context on generative ownership design.	general reference	Publisher abstract describes five generative-ownership patterns; the corpus's four principles remain its synthesis.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B025-O010	live	docs/humanized-autonomous-organization.md:947	kelly 2012	Generative ownership provides a broader tradition for the corpus's synthesized equity principles.	supported	Publisher abstract describes five generative-ownership patterns; the corpus's four principles remain its synthesis.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B025-O011	removed	docs/04-economics/04-equity-and-member-compensation.md:18	kelly 2012	The pre-audit sentence attributed the corpus's four compensation and ownership principles to stakeholder-alignment frameworks associated with Kelly.	partially supported	The publisher abstract describes five generative-ownership patterns, but not the corpus's four named compensation principles.	Attribute only the broader generative-ownership tradition to Kelly and identify the four principles as corpus synthesis.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B025-O012	removed	docs/humanized-autonomous-organization.md:949	kelly 2012	Consolidated pre-audit duplicate attributing the corpus's four principles to stakeholder-alignment frameworks associated with Kelly.	partially supported	The publisher abstract describes five generative-ownership patterns, but not the corpus's four named compensation principles.	Attribute only the broader generative-ownership tradition to Kelly and identify the four principles as corpus synthesis.
B025-O013	removed	docs/04-economics/03-sep-economic-agreements.md:96	kelly 2012	The pre-audit reference used the incomplete title <i>Owning Our Future: The Emerging Ownership Revolution</i> .	general reference	The publisher record gives the full title ending <i>Journeys to a Generative Economy</i> .	Supersede the incomplete title with the canonical citation.
B025-O014	removed	docs/04-economics/04-equity-and-member-compensation.md:119	kelly 2012	The pre-audit reference used only the short title <i>Owning Our Future</i> .	general reference	The publisher record gives the complete title and 2012 edition metadata.	Supersede the incomplete title with the canonical citation.
B025-O015	removed	docs/04-economics/05-reinvestment-mechanisms.md:151	kelly 2012	The pre-audit reference used only the short title <i>Owning Our Future</i> .	general reference	The publisher record gives the complete title and 2012 edition metadata.	Supersede the incomplete title with the canonical citation.
B025-O016	removed	docs/04-economics/06-external-interface-economics.md:126	kelly 2012	The pre-audit reference used only the short title <i>Owning Our Future</i> .	general reference	The publisher record gives the complete title and 2012 edition metadata.	Supersede the incomplete title with the canonical citation.
B025-O017	removed	docs/humanized-autonomous-organization.md:930	kelly 2012	The consolidated pre-audit reference omitted the final subtitle segment.	general reference	The publisher record gives the full title ending <i>Journeys to a Generative Economy</i> .	Supersede the incomplete title with the canonical citation.
B025-O018	removed	docs/humanized-autonomous-organization.md:1050	kelly 2012	The consolidated pre-audit reference used only the short title <i>Owning Our Future</i> .	general reference	The publisher record gives the complete title and 2012 edition metadata.	Supersede the incomplete title with the canonical citation.
B025-O019	removed	docs/humanized-autonomous-organization.md:1204	kelly 2012	The consolidated pre-audit reference used only the short title <i>Owning Our Future</i> .	general reference	The publisher record gives the complete title and 2012 edition metadata.	Supersede the incomplete title with the canonical citation.
B025-O020	removed	docs/humanized-autonomous-organization.md:1334	kelly 2012	The consolidated pre-audit reference used only the short title <i>Owning Our Future</i> .	general reference	The publisher record gives the complete title and 2012 edition metadata.	Supersede the incomplete title with the canonical citation.

B026 — Code: Version 2.0

- Canonical citation: Lessig, L. (2006). *Code: Version 2.0* (2nd ed.). Basic Books.
- Metadata status: verified; pre-audit metadata repaired
- Primary record: [Lessig's official book page](#)
- Disposition: retained
- Defect tags: code-version-title
- Forbidden live strings: Code: And Other Laws of Cyberspace, Version 2.0

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B026-O001	live	docs/13-member-trust-union/07-ethical-and-regulatory-anchoring.md:134	lessig 2006	Reference-list placement provides bounded context on code as a mode of regulation.	general reference	Author's official book page and section descriptions treat code as architectural regulation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B026-O002	live	docs/humanized-autonomous-organization.md:6076	lessig 2006	Reference-list placement provides bounded context on code as a mode of regulation.	general reference	Author's official book page and section descriptions treat code as architectural regulation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B026-O003	removed	docs/13-member-trust-union/07-ethical-and-regulatory-anchoring.md:134	lessig 2006	The pre-audit reference conflated the 1999 title with the revised work as <i>Code: And Other Laws of Cyberspace, Version 2.0</i> .	general reference	Lessig's official book page and the law-library catalogue identify the 2006 second edition as <i>Code: Version 2.0</i> .	Correct the conflated title while retaining the work as scoped regulatory context.
B026-O004	removed	docs/humanized-autonomous-organization.md:6090	lessig 2006	Consolidated pre-audit duplicate used the conflated 1999/2006 title.	general reference	Lessig's official book page and the law-library catalogue identify the 2006 second edition as <i>Code: Version 2.0</i> .	Correct the conflated title while retaining the work as scoped regulatory context.

B027 — Money and Sustainability: The Missing Link

- Canonical citation: Lietaer, B., Arnsperger, C., Goerner, S., & Brunnhuber, S. (2012). *Money and Sustainability: The Missing Link*. Triarchy Press.
- Metadata status: verified against the linked primary record
- Primary record: [Triarchy Press book record](#)
- Disposition: retained
- Defect tags: none
- Forbidden live strings: none

Occurrence State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B027-O001 live	docs/13-member-trust-union/03-governance-and-risk-distribution.md:157	lietaer arnsperger goerner brunnhuber 2012	Reference-list placement provides bounded context on monetary diversity and resilience.	general reference	Official executive summary section on monetary ecosystems argues for complementary monetary diversity.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B027-O002 live	docs/13-member-trust-union/06-integration-with-icn-and-external-markets.md:143	lietaer arnsperger goerner brunnhuber 2012	Reference-list placement provides bounded context on monetary diversity and resilience.	general reference	Official executive summary section on monetary ecosystems argues for complementary monetary diversity.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B027-O003 live	docs/13-member-trust-union/08-mtu-as-domain-specific-hao-implementation.md:113	lietaer arnsperger goerner brunnhuber 2012	Reference-list placement provides bounded context on monetary diversity and resilience.	general reference	Official executive summary section on monetary ecosystems argues for complementary monetary diversity.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B027-O004 live	docs/humanized-autonomous-organization.md:5518	lietaer arnsperger goerner brunnhuber 2012	Reference-list placement provides bounded context on monetary diversity and resilience.	general reference	Official executive summary section on monetary ecosystems argues for complementary monetary diversity.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B027-O005 live	docs/humanized-autonomous-organization.md:5936	lietaer arnsperger goerner brunnhuber 2012	Reference-list placement provides bounded context on monetary diversity and resilience.	general reference	Official executive summary section on monetary ecosystems argues for complementary monetary diversity.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B027-O006 live	docs/humanized-autonomous-organization.md:6200	lietaer arnsperger goerner brunnhuber 2012	Reference-list placement provides bounded context on monetary diversity and resilience.	general reference	Official executive summary section on monetary ecosystems argues for complementary monetary diversity.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

B028 — Interfirm Relationships and Informal Credit in Vietnam

- **Canonical citation:** McMillan, J., & Woodruff, C. (1999). Interfirm relationships and informal credit in Vietnam. *The Quarterly Journal of Economics*, 114(4), 1285–1320. <https://doi.org/10.1162/003355399556278>
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** [Oxford Academic article record](#)
- **Disposition:** retained
- **Defect tags:** informal-credit-doi
- **Forbidden live strings:** 10.1162/003355399556307

Occurrence State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B028-O001 live	docs/13-member-trust-union/01-multi-level-trust-architecture.md:113	mcmillan woodruff 1999	Reference-list placement provides bounded context on relational contracting and network-supported informal credit.	general reference	Official journal abstract links supplier credit to relationship duration, information, and business networks.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B028-O002 live	docs/humanized-autonomous-organization.md:5194	mcmillan woodruff 1999	Reference-list placement provides bounded context on relational contracting and network-supported informal credit.	general reference	Official journal abstract links supplier credit to relationship duration, information, and business networks.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B028-O003 removed	docs/13-member-trust-union/01-multi-level-trust-architecture.md:119	mcmillan woodruff 1999	The pre-audit reference gave DOI 10.1162/003355399556307 and omitted “The” from the journal title.	general reference	The Oxford Academic article record identifies <i>The Quarterly Journal of Economics</i> and DOI 10.1162/003355399556278.	Correct the DOI and journal title while retaining the work as scoped context.
B028-O004 removed	docs/humanized-autonomous-organization.md:5200	mcmillan woodruff 1999	Consolidated pre-audit duplicate carried the wrong DOI and shortened journal title.	general reference	The Oxford Academic article record identifies <i>The Quarterly Journal of Economics</i> and DOI 10.1162/003355399556278.	Correct the DOI and journal title while retaining the work as scoped context.

B029 — Bitcoin and Cryptocurrency Technologies

- **Canonical citation:** Narayanan, A., Bonneau, J., Felten, E., Miller, A., & Goldfeder, S. (2016). *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction*. Princeton University Press.
- **Metadata status:** verified against the linked primary record
- **Primary record:** [Authors' official Princeton-hosted site](#)
- **Disposition:** retained
- **Defect tags:** none
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B029-O001	live	docs/13-member-trust-union/01-multi-level-trust-architecture.md:105	narayanan bonneau felten miller goldfeder 2016	Reference-list placement provides bounded context on cryptographic identity, shared control, and reputation-system limits.	general reference	Official manuscript §1.4, pp. 41–42, and pp. 299–300 cover decentralized identity and reputation limits.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B029-O002	live	docs/13-member-trust-union/01-multi-level-trust-architecture.md:21	narayanan et al 2016	Cryptographic identity, shared control, and reputation-system limits inform a corpus-authored trust design.	supported	Official manuscript §1.4, pp. 41–42, and pp. 299–300 cover decentralized identity and reputation limits.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B029-O003	live	docs/humanized-autonomous-organization.md:5096	narayanan et al 2016	Cryptographic identity, shared control, and reputation-system limits inform a corpus-authored trust design.	supported	Official manuscript §1.4, pp. 41–42, and pp. 299–300 cover decentralized identity and reputation limits.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B029-O004	live	docs/humanized-autonomous-organization.md:5186	narayanan bonneau felten miller goldfeder 2016	Reference-list placement provides bounded context on cryptographic identity, shared control, and reputation-system limits.	general reference	Official manuscript §1.4, pp. 41–42, and pp. 299–300 cover decentralized identity and reputation limits.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B029-O005	removed	docs/13-member-trust-union/01-multi-level-trust-architecture.md:21	narayanan et al 2016	The pre-audit sentence attributed relational finance, reputation systems, multi-factor identity, and an alternative to FICO scoring to Narayanan and Ostrom.	partially supported	The official manuscript §1.4, pp. 41–42, and pp. 299–300 establish cryptographic identity and reputation limits; it does not address relational finance, FICO, or multi-factor identity.	Narrow the attribution to cryptographic identity, shared control, and reputation limits; identify the staged credit model as corpus design.
B029-O006	removed	docs/humanized-autonomous-organization.md:5096	narayanan et al 2016	Consolidated pre-audit duplicate attributing the relational-credit and FICO-alternative design to Narayanan and Ostrom.	partially supported	The official manuscript §1.4, pp. 41–42, and pp. 299–300 establish cryptographic identity and reputation limits; it does not address relational finance, FICO, or multi-factor identity.	Narrow the attribution to cryptographic identity, shared control, and reputation limits; identify the staged credit model as corpus design.

B030 — The Design of Everyday Things

- **Canonical citation:** Norman, D. A. (2013). *The Design of Everyday Things: Revised and Expanded Edition*. Basic Books.
- **Metadata status:** verified; pre-audit revised–edition title completed
- **Primary record:** [Basic Books/Hachette edition record](#)
- **Disposition:** retained
- **Defect tags:** design-everyday-things-title
- **Forbidden live strings:** Norman, D. A. (2013). *The Design of Everyday Things* (Revised edition). Basic Books.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B030-O001	live	docs/07-human-systems/00-overview.md:29	norman 2013	Reference-list placement provides bounded context on human-centred design.	general reference	Publisher abstract identifies the book as a guide to human-centred design, user needs, constraints, and feedback.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B030-O002	live	docs/07-human-systems/00-overview.md:9	norman 2013	Human-centred design is one input to the human-systems synthesis.	supported	Publisher abstract identifies the book as a guide to human-centred design, user needs, constraints, and feedback.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B030-O003	live	docs/13-member-trust-union/04-technology-and-physical-digital-integration.md:131	norman 2013	Reference-list placement provides bounded context on human-centred design.	general reference	Publisher abstract identifies the book as a guide to human-centred design, user needs, constraints, and feedback.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B030-O004	live	docs/humanized-autonomous-organization.md:1977	norman 2013	Human-centred design is one input to the human-systems synthesis.	supported	Publisher abstract identifies the book as a guide to human-centred design, user needs, constraints, and feedback.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B030-O005	live	docs/humanized-autonomous-organization.md:1997	norman 2013	Reference-list placement provides bounded context on human-centred design.	general reference	Publisher abstract identifies the book as a guide to human-centred design, user needs, constraints, and feedback.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B030-O006	live	docs/humanized-autonomous-organization.md:5651	norman 2013	Reference-list placement provides bounded context on human-centred design.	general reference	Publisher abstract identifies the book as a guide to human-centred design, user needs, constraints, and feedback.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B030-O007	removed	docs/07-human-systems/00-overview.md:9	norman 2013	The pre-audit sentence attached human-centered design to intended HAO resilience and inclusion as the organization grows.	partially supported	The publisher abstract identifies human-centered design, user needs, constraints, and feedback, but not the claimed HAO outcomes or scaling effects.	Preserve the supported design influence and identify the HAO outcomes as corpus intentions.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B030-O008	removed	docs/humanized-autonomous-organization.md:1983	norman 2013	The consolidated pre-audit sentence said the cited disciplines ensure HAO resilience, inclusion, and meaning generation as the organization grows.	partially supported	The publisher abstract identifies human-centered design, user needs, constraints, and feedback, but not guaranteed HAO outcomes or scaling effects.	Replace the guarantee with a bounded influence statement and corpus-owned design intention.
B030-O009	removed	docs/13-member-trust-union/04-technology-and-physical-digital-integration.md:131	norman 2013	The pre-audit reference used <i>The Design of Everyday Things</i> with a generic revised-edition note.	general reference	The Basic Books edition record gives <i>The Design of Everyday Things: Revised and Expanded Edition</i> .	Supersede the shortened edition title with the canonical citation.
B030-O010	removed	docs/humanized-autonomous-organization.md:5663	norman 2013	Consolidated pre-audit duplicate used the shortened title and generic revised-edition note.	general reference	The Basic Books edition record gives <i>The Design of Everyday Things: Revised and Expanded Edition</i> .	Supersede the shortened edition title with the canonical citation.

B031 — Social Lending: Transforming Traditional Credit Systems

- **Canonical citation:** No canonical Nyer and Smith (2013) article was established; the cited pages are occupied by a different Wiley article.
- **Metadata status:** baseline identity contradicted by official DOI or issue metadata
- **Primary record:** Official DOI record for the article occupying the cited pages
- **Disposition:** removed
- **Defect tags:** nyer-social-lending
- **Forbidden live strings:** Social Lending: Transforming Traditional Credit Systems

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B031-O001	removed	docs/13-member-trust-union/02-financial-product-design.md:165	nyer smith 2013	Historical pre-audit occurrence using the superseded or removed nyer smith 2013 attribution.	unsupported	The official DOI and Wiley issue metadata assign the claimed pages to Key and coauthors' different article.	Remove the citation without substituting the occupying article.
B031-O002	removed	docs/humanized-autonomous-organization.md:5371	nyer smith 2013	Historical pre-audit occurrence using the superseded or removed nyer smith 2013 attribution.	unsupported	The official DOI and Wiley issue metadata assign the claimed pages to Key and coauthors' different article.	Remove the citation without substituting the occupying article.

B032 — Governing the Commons

- **Canonical citation:** Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511807763>
- **Metadata status:** verified against the linked primary record
- **Primary record:** Cambridge University Press book record
- **Disposition:** retained
- **Defect tags:** none
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B032-O001	live	docs/13-member-trust-union/02-financial-product-design.md:159	ostrom 1990	Reference-list placement provides bounded context on self-organized common-pool-resource governance.	general reference	Cambridge synopsis and chapter 2 establish self-organization and governance of common-pool resources.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B032-O002	live	docs/13-member-trust-union/08-mtu-as-domain-specific-hao-implementation.md:115	ostrom 1990	Reference-list placement provides bounded context on self-organized common-pool-resource governance.	general reference	Cambridge synopsis and chapter 2 establish self-organization and governance of common-pool resources.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B032-O003	live	docs/humanized-autonomous-organization.md:5359	ostrom 1990	Reference-list placement provides bounded context on self-organized common-pool-resource governance.	general reference	Cambridge synopsis and chapter 2 establish self-organization and governance of common-pool resources.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B032-O004	live	docs/humanized-autonomous-organization.md:6202	ostrom 1990	Reference-list placement provides bounded context on self-organized common-pool-resource governance.	general reference	Cambridge synopsis and chapter 2 establish self-organization and governance of common-pool resources.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

B033 — Understanding Institutional Diversity

- **Canonical citation:** Ostrom, E. (2005). *Understanding Institutional Diversity*. Princeton University Press.
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** Official JSTOR/Princeton record
- **Disposition:** retained

- Defect tags: `ostrom-understanding-year`
- Forbidden live strings: `Ostrom, E. (2009). *Understanding Institutional Diversity*`

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B033-O001	live	<code>docs/13-member-trust-union/03-governance-and-risk-distribution.md:149</code>	<code>ostrom 2005</code>	Reference-list placement provides bounded context on institutional analysis and rule systems.	<code>general reference</code>	Official publisher/JSTOR contents establish institutional rules and IAD subject matter, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B033-O002	live	<code>docs/13-member-trust-union/06-integration-with-icn-and-external-markets.md:147</code>	<code>ostrom 2005</code>	Reference-list placement provides bounded context on institutional analysis and rule systems.	<code>general reference</code>	Official publisher/JSTOR contents establish institutional rules and IAD subject matter, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B033-O003	live	<code>docs/13-member-trust-union/07-ethical-and-regulatory-anchoring.md:136</code>	<code>ostrom 2005</code>	Reference-list placement provides bounded context on institutional analysis and rule systems.	<code>general reference</code>	Official publisher/JSTOR contents establish institutional rules and IAD subject matter, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B033-O004	live	<code>docs/humanized-autonomous-organization.md:5510</code>	<code>ostrom 2005</code>	Reference-list placement provides bounded context on institutional analysis and rule systems.	<code>general reference</code>	Official publisher/JSTOR contents establish institutional rules and IAD subject matter, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B033-O005	live	<code>docs/humanized-autonomous-organization.md:5940</code>	<code>ostrom 2005</code>	Reference-list placement provides bounded context on institutional analysis and rule systems.	<code>general reference</code>	Official publisher/JSTOR contents establish institutional rules and IAD subject matter, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B033-O006	live	<code>docs/humanized-autonomous-organization.md:6078</code>	<code>ostrom 2005</code>	Reference-list placement provides bounded context on institutional analysis and rule systems.	<code>general reference</code>	Official publisher/JSTOR contents establish institutional rules and IAD subject matter, not MTU mechanisms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B033-O007	removed	<code>docs/13-member-trust-union/03-governance-and-risk-distribution.md:149</code>	<code>ostrom 2009</code>	The pre-audit reference dated <i>Understanding Institutional Diversity</i> 2009 rather than 2005.	<code>general reference</code>	The Princeton/JSTOR record identifies the book as a 2005 publication.	Correct the year while retaining the book as scoped institutional-analysis context.
B033-O008	removed	<code>docs/humanized-autonomous-organization.md:5520</code>	<code>ostrom 2009</code>	The consolidated pre-audit reference dated <i>Understanding Institutional Diversity</i> 2009 rather than 2005.	<code>general reference</code>	The Princeton/JSTOR record identifies the book as a 2005 publication.	Correct the year while retaining the book as scoped institutional-analysis context.

B034 — Beyond Markets and States

- Canonical citation: Ostrom, E. (2010). Beyond markets and states: Polycentric governance of complex economic systems. *American Economic Review*, 100(3), 641–672. <https://doi.org/10.1257/aer.100.3.641>
- Metadata status: `verified; eight pre-audit reference-title variants completed`
- Primary record: [American Economic Association article record](#)
- Disposition: `retained`
- Defect tags: `beyond-markets-title`
- Forbidden live strings: `Ostrom, E. (2010). *Beyond Markets and States*`

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B034-O001	live	<code>docs/04-economics/00-overview.md:8</code>	<code>ostrom 2010</code>	Ostrom's polycentric coordination or trust-and-reputation premise is applied as a corpus design.	<code>supported</code>	Official Nobel full text §II.A and §§V–VI supports polycentric coordination and trust/reputation in cooperation.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B034-O002	live	<code>docs/04-economics/01-trickle-out-investment-flow.md:87</code>	<code>ostrom 2010</code>	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	<code>general reference</code>	Official Nobel full text §II.A and §§V–VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O003	live	<code>docs/04-economics/03-sep-economic-agreements.md:94</code>	<code>ostrom 2010</code>	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	<code>general reference</code>	Official Nobel full text §II.A and §§V–VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O004	live	<code>docs/04-economics/04-equity-and-member-compensation.md:118</code>	<code>ostrom 2010</code>	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	<code>general reference</code>	Official Nobel full text §II.A and §§V–VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O005	live	<code>docs/04-economics/05-reinvestment-mechanisms.md:149</code>	<code>ostrom 2010</code>	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	<code>general reference</code>	Official Nobel full text §II.A and §§V–VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O006	live	<code>docs/04-economics/06-external-interface-economics.md:127</code>	<code>ostrom 2010</code>	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	<code>general reference</code>	Official Nobel full text §II.A and §§V–VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B034-O007	live	docs/12-icn/02-governance-framework-and-dea-protocols.md:12	ostrom 2010	Ostrom's polycentric coordination or trust-and-reputation premise is applied as a corpus design.	supported	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B034-O008	live	docs/13-member-trust-union/01-multi-level-trust-architecture.md:107	ostrom 2010	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	general reference	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O009	live	docs/13-member-trust-union/01-multi-level-trust-architecture.md:21	ostrom 2010	Ostrom's polycentric coordination or trust-and-reputation premise is applied as a corpus design.	supported	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B034-O010	live	docs/13-member-trust-union/04-technology-and-physical-digital-integration.md:133	ostrom 2010	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	general reference	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O011	live	docs/humanized-autonomous-organization.md:1047	ostrom 2010	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	general reference	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O012	live	docs/humanized-autonomous-organization.md:1199	ostrom 2010	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	general reference	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O013	live	docs/humanized-autonomous-organization.md:1330	ostrom 2010	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	general reference	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O014	live	docs/humanized-autonomous-organization.md:4424	ostrom 2010	Ostrom's polycentric coordination or trust-and-reputation premise is applied as a corpus design.	supported	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B034-O015	live	docs/humanized-autonomous-organization.md:5096	ostrom 2010	Ostrom's polycentric coordination or trust-and-reputation premise is applied as a corpus design.	supported	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B034-O016	live	docs/humanized-autonomous-organization.md:5188	ostrom 2010	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	general reference	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O017	live	docs/humanized-autonomous-organization.md:5653	ostrom 2010	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	general reference	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O018	live	docs/humanized-autonomous-organization.md:636	ostrom 2010	Ostrom's polycentric coordination or trust-and-reputation premise is applied as a corpus design.	supported	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B034-O019	live	docs/humanized-autonomous-organization.md:735	ostrom 2010	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	general reference	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O020	live	docs/humanized-autonomous-organization.md:927	ostrom 2010	Reference-list placement provides bounded context on polycentric coordination and trust in collective action.	general reference	Official Nobel full text §II.A and §§V-VI supports polycentric coordination and trust/reputation in cooperation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B034-O021	removed	docs/04-economics/00-overview.md:8	ostrom 2010	The pre-audit sentence said the whole economic model draws on cooperative economics, attaching Ostrom to the ensuing HAO finance mechanics.	partially supported	The official Nobel full text §II.A supports polycentric institutional coordination beyond a market/state binary, not the corpus's capital-allocation mechanics.	Limit the attribution to the governance premise and identify the economic mechanics as corpus proposals.
B034-O022	removed	docs/12-icn/02-governance-framework-and-dea-protocols.md:12	ostrom 2010	The pre-audit sentence said AGF draws on polycentric theory to balance autonomy and network-wide coherence.	partially supported	The official Nobel full text §II.A describes multiple formally independent decision centers capable of coherent coordination, but not AGF, DEA, sociocracy, or design lineage.	State the supported premise and identify AGF as the corpus's application.
B034-O023	removed	docs/13-member-trust-union/01-multi-level-trust-architecture.md:21	ostrom 2010	The pre-audit sentence attributed relational finance, reputation systems, multi-factor identity, and an alternative to FICO scoring to Ostrom and Narayanan.	partially supported	The official Nobel full text §§V-VI supports trust and reputation in collective action, but not relational finance, identity design, credit scoring, or FICO.	Retain only the trust-and-reputation premise and identify the MTU credit design as a corpus proposal.
B034-O024	removed	docs/humanized-autonomous-organization.md:636	ostrom 2010	The consolidated pre-audit sentence said the economic model is guided by regenerative finance and cooperative economics.	partially supported	The official Nobel full text §II.A supports polycentric institutional coordination beyond a market/state binary, not regenerative finance or the corpus's capital-allocation mechanics.	Limit the attribution to the governance premise and identify the economic mechanics as corpus proposals.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B034-O025	removed	docs/humanized-autonomous-organization.md:4424	ostrom 2010	The consolidated pre-audit sentence described AGF as inspired by polycentric governance while bundling AGF-specific mechanisms.	partially supported	The official Nobel full text §II.A describes independent decision centers capable of coherent coordination, but not AGF, DEA, sociocracy, or historical design influence.	State the supported premise and identify AGF as the corpus's application.
B034-O026	removed	docs/humanized-autonomous-organization.md:5096	ostrom 2010	The consolidated pre-audit sentence attributed the relational-credit and FICO-alternative design to Ostrom and Narayanan.	partially supported	The official Nobel full text §§V–VI supports trust and reputation in collective action, but not relational finance, identity design, credit scoring, or FICO.	Retain only the trust-and-reputation premise and identify the MTU credit design as a corpus proposal.
B034-O027	removed	docs/04-economics/03-sep-economic-agreements.md:95	ostrom 2010	The pre-audit reference truncated the article title to <i>Beyond Markets and States</i> .	general reference	The AEA record identifies the complete title <i>Beyond Markets and States: Polycentric Governance of Complex Economic Systems</i> .	Complete the title while retaining the article as scoped polycentric-governance context.
B034-O028	removed	docs/humanized-autonomous-organization.md:929	ostrom 2010	Consolidated pre-audit duplicate truncated the article title to <i>Beyond Markets and States</i> .	general reference	The AEA record identifies the complete title <i>Beyond Markets and States: Polycentric Governance of Complex Economic Systems</i> .	Complete the title while retaining the article as scoped polycentric-governance context.
B034-O029	removed	docs/04-economics/04-equity-and-member-compensation.md:118	ostrom 2010	The pre-audit reference truncated the article title to <i>Beyond Markets and States</i> .	general reference	The AEA record identifies the complete title <i>Beyond Markets and States: Polycentric Governance of Complex Economic Systems</i> .	Complete the title while retaining the article as scoped polycentric-governance context.
B034-O030	removed	docs/04-economics/05-reinvestment-mechanisms.md:150	ostrom 2010	The pre-audit reference truncated the article title to <i>Beyond Markets and States</i> .	general reference	The AEA record identifies the complete title <i>Beyond Markets and States: Polycentric Governance of Complex Economic Systems</i> .	Complete the title while retaining the article as scoped polycentric-governance context.
B034-O031	removed	docs/04-economics/06-external-interface-economics.md:127	ostrom 2010	The pre-audit reference truncated the article title to <i>Beyond Markets and States</i> .	general reference	The AEA record identifies the complete title <i>Beyond Markets and States: Polycentric Governance of Complex Economic Systems</i> .	Complete the title while retaining the article as scoped polycentric-governance context.
B034-O032	removed	docs/humanized-autonomous-organization.md:1049	ostrom 2010	Consolidated pre-audit duplicate truncated the article title to <i>Beyond Markets and States</i> .	general reference	The AEA record identifies the complete title <i>Beyond Markets and States: Polycentric Governance of Complex Economic Systems</i> .	Complete the title while retaining the article as scoped polycentric-governance context.
B034-O033	removed	docs/humanized-autonomous-organization.md:1203	ostrom 2010	Consolidated pre-audit duplicate truncated the article title to <i>Beyond Markets and States</i> .	general reference	The AEA record identifies the complete title <i>Beyond Markets and States: Polycentric Governance of Complex Economic Systems</i> .	Complete the title while retaining the article as scoped polycentric-governance context.
B034-O034	removed	docs/humanized-autonomous-organization.md:1335	ostrom 2010	Consolidated pre-audit duplicate truncated the article title to <i>Beyond Markets and States</i> .	general reference	The AEA record identifies the complete title <i>Beyond Markets and States: Polycentric Governance of Complex Economic Systems</i> .	Complete the title while retaining the article as scoped polycentric-governance context.

B035 — Governing the Commons cited with the 2010 date

- **Canonical citation:** Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511807763>
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** [Cambridge front matter recording first publication in 1990](#)
- **Disposition:** merged
- **Defect tags:** ostrom-governing-title-year
- **Forbidden live strings:** Ostrom, E. (2010). *Governing the Commons.*

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B035-O001	live	docs/04-economics/02-revenue-allocation-framework.md:91	ostrom 1990	Reference-list placement provides bounded context on self-organized common-pool-resource governance.	general reference	Cambridge front matter states first publication in 1990; chapter 2 concerns self-governance.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B035-O002	live	docs/humanized-autonomous-organization.md:830	ostrom 1990	Reference-list placement provides bounded context on self-organized common-pool-resource governance.	general reference	Cambridge front matter states first publication in 1990; chapter 2 concerns self-governance.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B035-O003	removed	docs/04-economics/02-revenue-allocation-framework.md:91	ostrom 2010	Historical pre-audit occurrence using the superseded or removed ostrom 2010 attribution.	unsupported	Cambridge front matter, p. iv, states that <i>Governing the Commons</i> was first published in 1990.	Correct the year to 1990 and merge the repaired occurrence with B032's canonical work.
B035-O004	removed	docs/humanized-autonomous-organization.md:830	ostrom 2010	Historical pre-audit occurrence using the superseded or removed ostrom 2010 attribution.	unsupported	Cambridge front matter, p. iv, states that <i>Governing the Commons</i> was first published in 1990.	Correct the year to 1990 and merge the repaired occurrence with B032's canonical work.

B036 — Commons-Oriented Equity Structures

- **Canonical citation:** No canonical P2P Foundation (2020) publication was established for the baseline string.

- **Metadata status:** not corroborated; Bauwens and Niaros (2017) is a distinct candidate work, not the baseline P2P Foundation publication
- **Primary record:** No primary record for the baseline identity was located. Comparison only: [official P2P Foundation catalogue entry](#) and [Heinrich Böll Foundation co-publisher record](#) for the distinct Bauwens and Niaros (2017) report.
- **Disposition:** removed
- **Defect tags:** p2p-equity-structures
- **Forbidden live strings:** Commons-Oriented Equity Structures

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B036-O001	removed	docs/04-economics/04-equity-and-member-compensation.md:120	p2p foundation 2020	Pre-audit reference to the uncorroborated <i>Commons-Oriented Equity Structures</i> identity.	not verifiable	The search could not locate the baseline record. The distinct 2017 report is broadly relevant to contributory value accounting but not the chapter's equity instruments, vesting, portability, or recovery rules.	Remove the invalid entry. Do not introduce the candidate report in this repair; audit it separately if later added.
B036-O002	removed	docs/humanized-autonomous-organization.md:1051	p2p foundation 2020	Consolidated duplicate of the uncorroborated <i>Commons-Oriented Equity Structures</i> identity.	not verifiable	The search could not locate the baseline record. The distinct 2017 report is broadly relevant to contributory value accounting but not the chapter's equity instruments, vesting, portability, or recovery rules.	Remove the invalid entry. Do not introduce the candidate report in this repair; audit it separately if later added.

B037 — Building the New Economy: Data as Capital

- **Canonical citation:** Pentland, A., Lipton, A., & Hardjono, T. (2021). *Building the New Economy: Data as Capital*. MIT Press. <https://doi.org/10.7551/mitpress/13991.001.0001>
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** [MIT Press book record](#)
- **Disposition:** retained
- **Defect tags:** data-as-capital-authors
- **Forbidden live strings:** Wozniak, P. (2021) ; Hardjono, T., & Pentland, A. (2021)

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B037-O001	live	docs/13-member-trust-union/01-multi-level-trust-architecture.md:115	pentland lipton hardjono 2021	Reference-list placement provides bounded context on data trusts, accountability, and interoperability.	general reference	MIT Press abstract covers data trusts, accountable systems, secure transactions, and interoperability.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B037-O002	live	docs/13-member-trust-union/04-technology-and-physical-digital-integration.md:137	pentland lipton hardjono 2021	Reference-list placement provides bounded context on data trusts, accountability, and interoperability.	general reference	MIT Press abstract covers data trusts, accountable systems, secure transactions, and interoperability.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B037-O003	live	docs/13-member-trust-union/06-integration-with-icn-and-external-markets.md:149	pentland lipton hardjono 2021	Reference-list placement provides bounded context on data trusts, accountability, and interoperability.	general reference	MIT Press abstract covers data trusts, accountable systems, secure transactions, and interoperability.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B037-O004	live	docs/humanized-autonomous-organization.md:5196	pentland lipton hardjono 2021	Reference-list placement provides bounded context on data trusts, accountability, and interoperability.	general reference	MIT Press abstract covers data trusts, accountable systems, secure transactions, and interoperability.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B037-O005	live	docs/humanized-autonomous-organization.md:5657	pentland lipton hardjono 2021	Reference-list placement provides bounded context on data trusts, accountability, and interoperability.	general reference	MIT Press abstract covers data trusts, accountable systems, secure transactions, and interoperability.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B037-O006	live	docs/humanized-autonomous-organization.md:5942	pentland lipton hardjono 2021	Reference-list placement provides bounded context on data trusts, accountability, and interoperability.	general reference	MIT Press abstract covers data trusts, accountable systems, secure transactions, and interoperability.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B037-O007	removed	docs/13-member-trust-union/01-multi-level-trust-architecture.md:121	pentland hardjono wozniak 2021	The pre-audit reference omitted coauthor Alexander Lipton and substituted Paul Wozniak in the author list.	general reference	The MIT Press record identifies Alex Pentland, Alexander Lipton, and Thomas Hardjono as the book's authors.	Restore the verified author list while retaining the book as scoped context.
B037-O008	removed	docs/13-member-trust-union/04-technology-and-physical-digital-integration.md:137	pentland hardjono wozniak 2021	The pre-audit reference omitted coauthor Alexander Lipton and substituted Paul Wozniak in the author list.	general reference	The MIT Press record identifies Alex Pentland, Alexander Lipton, and Thomas Hardjono as the book's authors.	Restore the verified author list while retaining the book as scoped context.
B037-O009	removed	docs/13-member-trust-union/06-integration-with-icn-and-external-markets.md:149	hardjono pentland 2021	The pre-audit reference omitted coauthor Alexander Lipton and reversed the order of the two named authors.	general reference	The MIT Press record identifies Alex Pentland, Alexander Lipton, and Thomas Hardjono, in that order, as the book's authors.	Restore the verified author list and order while retaining the book as scoped context.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B037-O010	removed	docs/humanized-autonomous-organization.md:5202	pentland hardjono wozniak 2021	The consolidated pre-audit reference omitted coauthor Alexander Lipton and substituted Paul Wozniak in the author list.	general reference	The MIT Press record identifies Alex Pentland, Alexander Lipton, and Thomas Hardjono as the book's authors.	Restore the verified author list while retaining the book as scoped context.
B037-O011	removed	docs/humanized-autonomous-organization.md:5669	pentland hardjono wozniak 2021	The consolidated pre-audit reference omitted coauthor Alexander Lipton and substituted Paul Wozniak in the author list.	general reference	The MIT Press record identifies Alex Pentland, Alexander Lipton, and Thomas Hardjono as the book's authors.	Restore the verified author list while retaining the book as scoped context.
B037-O012	removed	docs/humanized-autonomous-organization.md:5954	hardjono pentland 2021	The consolidated pre-audit reference omitted coauthor Alexander Lipton and reversed the order of the two named authors.	general reference	The MIT Press record identifies Alex Pentland, Alexander Lipton, and Thomas Hardjono, in that order, as the book's authors.	Restore the verified author list and order while retaining the book as scoped context.

B038 — Capital in the Twenty-First Century

- **Canonical citation:** Piketty, T. (2014). *Capital in the Twenty-First Century* (A. Goldhammer, Trans.). The Belknap Press of Harvard University Press.
- **Metadata status:** verified against the 2014 edition catalogue; the author-associated WID text is the claim source, not the edition-metadata record
- **Primary record:** Authoritative 2014 edition catalogue record; author-associated introduction and chapter 1
- **Disposition:** retained
- **Defect tags:** none
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B038-O001	live	docs/13-member-trust-union/05-member-experience-and-compensation.md:126	piketty 2014	Reference-list placement provides bounded context on wealth and income concentration.	general reference	Author-associated introduction, chapter 1, frames historical income and wealth concentration.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B038-O002	live	docs/humanized-autonomous-organization.md:5785	piketty 2014	Reference-list placement provides bounded context on wealth and income concentration.	general reference	Author-associated introduction, chapter 1, frames historical income and wealth concentration.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

B039 — Working Together

- **Canonical citation:** Poteete, A. R., Janssen, M. A., & Ostrom, E. (2010). *Working Together: Collective Action, the Commons, and Multiple Methods in Practice*. Princeton University Press. <https://doi.org/10.1515/9781400835157>
- **Metadata status:** verified against the linked primary record
- **Primary record:** Publisher DOI record
- **Disposition:** retained
- **Defect tags:** none
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B039-O001	live	docs/13-member-trust-union/03-governance-and-risk-distribution.md:159	poteete janssen ostrom 2010	Reference-list placement provides bounded context on collective action and commons research.	general reference	Publisher abstract describes multimethod collective-action and commons research, not MTU risk algorithms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B039-O002	live	docs/humanized-autonomous-organization.md:5520	poteete janssen ostrom 2010	Reference-list placement provides bounded context on collective action and commons research.	general reference	Publisher abstract describes multimethod collective-action and commons research, not MTU risk algorithms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

B040 — Doughnut Economics

- **Canonical citation:** Raworth, K. (2017). *Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist*. Random House Business Books.
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** Penguin 2017 edition record
- **Disposition:** retained
- **Defect tags:** doughnut-edition-hybrid
- **Forbidden live strings:** ****Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist**** – Raworth, K. (2017). ***Chelsea Green.***; Raworth, K. (2017). ***Doughnut Economics***; Raworth, K. (2017). ***Doughnut Economics.***

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B040-O001	live	docs/04-economics/02-revenue-allocation-framework.md:93	raworth 2017	Reference-list placement provides bounded context on distributive and regenerative economic framing.	general reference	Publisher abstract describes distributive and regenerative design, not the corpus's allocation formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B040-O002	live	docs/04-economics/04-equity-and-member-compensation.md:120	raworth 2017	Reference-list placement provides bounded context on distributive and regenerative economic framing.	general reference	Publisher abstract describes distributive and regenerative design, not the corpus's allocation formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B040-O003	live	docs/04-economics/05-reinvest-mechanisms.md:148	raworth 2017	Reference-list placement provides bounded context on distributive and regenerative economic framing.	general reference	Publisher abstract describes distributive and regenerative design, not the corpus's allocation formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B040-O004	live	docs/04-economics/06-external-interface-economics.md:130	raworth 2017	Reference-list placement provides bounded context on distributive and regenerative economic framing.	general reference	Publisher abstract describes distributive and regenerative design, not the corpus's allocation formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B040-O005	live	docs/humanized-autonomous-organization.md:1049	raworth 2017	Reference-list placement provides bounded context on distributive and regenerative economic framing.	general reference	Publisher abstract describes distributive and regenerative design, not the corpus's allocation formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B040-O006	live	docs/humanized-autonomous-organization.md:1198	raworth 2017	Reference-list placement provides bounded context on distributive and regenerative economic framing.	general reference	Publisher abstract describes distributive and regenerative design, not the corpus's allocation formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B040-O007	live	docs/humanized-autonomous-organization.md:1333	raworth 2017	Reference-list placement provides bounded context on distributive and regenerative economic framing.	general reference	Publisher abstract describes distributive and regenerative design, not the corpus's allocation formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B040-O008	live	docs/humanized-autonomous-organization.md:832	raworth 2017	Reference-list placement provides bounded context on distributive and regenerative economic framing.	general reference	Publisher abstract describes distributive and regenerative design, not the corpus's allocation formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B040-O009	removed	docs/04-economics/02-revenue-allocation-framework.md:93	raworth 2017	The pre-audit reference used the short title <i>Doughnut Economics</i> .	general reference	Penguin's 2017 edition record identifies the complete title <i>Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist</i> .	Complete the title while retaining the work as scoped context.
B040-O010	removed	docs/04-economics/04-equity-and-member-compensation.md:121	raworth 2017	The pre-audit reference used the short title <i>Doughnut Economics</i> .	general reference	Penguin's 2017 edition record identifies the complete title <i>Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist</i> .	Complete the title while retaining the work as scoped context.
B040-O011	removed	docs/04-economics/05-reinvest-mechanisms.md:148	raworth 2017	The pre-audit reference used the short title <i>Doughnut Economics</i> .	general reference	Penguin's 2017 edition record identifies the complete title <i>Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist</i> .	Complete the title while retaining the work as scoped context.
B040-O012	removed	docs/04-economics/06-external-interface-economics.md:130	raworth 2017	The pre-audit reference used the short title <i>Doughnut Economics</i> .	general reference	Penguin's 2017 edition record identifies the complete title <i>Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist</i> .	Complete the title while retaining the work as scoped context.
B040-O013	removed	docs/humanized-autonomous-organization.md:832	raworth 2017	The consolidated pre-audit reference used the short title <i>Doughnut Economics</i> .	general reference	Penguin's 2017 edition record identifies the complete title <i>Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist</i> .	Complete the title while retaining the work as scoped context.
B040-O014	removed	docs/humanized-autonomous-organization.md:1052	raworth 2017	The consolidated pre-audit reference used the short title <i>Doughnut Economics</i> .	general reference	Penguin's 2017 edition record identifies the complete title <i>Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist</i> .	Complete the title while retaining the work as scoped context.
B040-O015	removed	docs/humanized-autonomous-organization.md:1201	raworth 2017	The consolidated pre-audit reference used the short title <i>Doughnut Economics</i> .	general reference	Penguin's 2017 edition record identifies the complete title <i>Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist</i> .	Complete the title while retaining the work as scoped context.
B040-O016	removed	docs/humanized-autonomous-organization.md:1338	raworth 2017	The consolidated pre-audit reference used the short title <i>Doughnut Economics</i> .	general reference	Penguin's 2017 edition record identifies the complete title <i>Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist</i> .	Complete the title while retaining the work as scoped context.

B041 — The Lean Startup

- **Canonical citation:** Ries, E. (2011). *The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*. Crown Business.
- **Metadata status:** verified against the linked primary record
- **Primary record:** [Penguin Random House publisher record](#)
- **Disposition:** retained
- **Defect tags:** none

- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B041-O001	live	docs/04-economics/01-trickle-out-investment-flow.md:13	ries 2011	Validated-learning cycles can inform assessment under a corpus-authored capital-release rule.	supported	Publisher abstract explicitly identifies validated learning and rapid scientific experimentation.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B041-O002	live	docs/04-economics/01-trickle-out-investment-flow.md:89	ries 2011	Reference-list placement provides bounded context on validated learning and rapid experimentation.	general reference	Publisher abstract explicitly identifies validated learning and rapid scientific experimentation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B041-O003	live	docs/humanized-autonomous-organization.md:661	ries 2011	Validated-learning cycles can inform assessment under a corpus-authored capital-release rule.	supported	Publisher abstract explicitly identifies validated learning and rapid scientific experimentation.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B041-O004	live	docs/humanized-autonomous-organization.md:737	ries 2011	Reference-list placement provides bounded context on validated learning and rapid experimentation.	general reference	Publisher abstract explicitly identifies validated learning and rapid scientific experimentation.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B041-O005	removed	docs/04-economics/01-trickle-out-investment-flow.md:13	ries 2011	The pre-audit sentence implied that milestones or validated-learning cycles unlock additional capital rounds.	partially supported	The publisher abstract establishes validated learning and rapid scientific experimentation, but not milestone-triggered release of later capital rounds.	Keep validated learning as an assessment input and identify capital release as a corpus rule.
B041-O006	removed	docs/humanized-autonomous-organization.md:661	ries 2011	Consolidated pre-audit duplicate implying that milestones or validated-learning cycles unlock additional capital rounds.	partially supported	The publisher abstract establishes validated learning and rapid scientific experimentation, but not milestone-triggered release of later capital rounds.	Keep validated learning as an assessment input and identify capital release as a corpus rule.

B042 — Platform Cooperativism

- **Canonical citation:** Scholz, T. (2016). *Platform Cooperativism: Challenging the Corporate Sharing Economy*. Rosa Luxemburg Stiftung—New York Office.
- **Metadata status:** verified; pre-audit title and publisher completed
- **Primary record:** [Rosa Luxemburg Stiftung publication record](#)
- **Disposition:** retained
- **Defect tags:** platform-cooperativism-title
- **Forbidden live strings:** Scholz, T. (2016). *Platform Cooperativism*

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B042-O001	live	docs/04-economics/03-sep-economic-agreements.md:93	scholz 2016	Reference-list placement provides bounded context on cooperative platform ownership.	general reference	Official report p. 14 advocates cooperative ownership and democratic governance, not SEP formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B042-O002	live	docs/04-economics/06-external-interface-economics.md:129	scholz 2016	Reference-list placement provides bounded context on cooperative platform ownership.	general reference	Official report p. 14 advocates cooperative ownership and democratic governance, not SEP formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B042-O003	live	docs/humanized-autonomous-organization.md:1332	scholz 2016	Reference-list placement provides bounded context on cooperative platform ownership.	general reference	Official report p. 14 advocates cooperative ownership and democratic governance, not SEP formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B042-O004	live	docs/humanized-autonomous-organization.md:926	scholz 2016	Reference-list placement provides bounded context on cooperative platform ownership.	general reference	Official report p. 14 advocates cooperative ownership and democratic governance, not SEP formulas.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B042-O005	removed	docs/04-economics/03-sep-economic-agreements.md:94	scholz 2016	The pre-audit reference used only the title <i>Platform Cooperativism</i> and omitted the publisher.	general reference	The official report, p. 14, and publication record identify <i>Platform Cooperativism: Challenging the Corporate Sharing Economy</i> and Rosa Luxemburg Stiftung—New York Office.	Complete the title and publisher while retaining the work as scoped context.
B042-O006	removed	docs/04-economics/06-external-interface-economics.md:129	scholz 2016	The pre-audit reference used only the title <i>Platform Cooperativism</i> and omitted the publisher.	general reference	The official report, p. 14, and publication record identify <i>Platform Cooperativism: Challenging the Corporate Sharing Economy</i> and Rosa Luxemburg Stiftung—New York Office.	Complete the title and publisher while retaining the work as scoped context.
B042-O007	removed	docs/humanized-autonomous-organization.md:928	scholz 2016	Consolidated pre-audit reference used only the title <i>Platform Cooperativism</i> and omitted the publisher.	general reference	The official report, p. 14, and publication record identify <i>Platform Cooperativism: Challenging the Corporate Sharing Economy</i> and Rosa Luxemburg Stiftung—New York Office.	Complete the title and publisher while retaining the work as scoped context.
B042-O008	removed	docs/humanized-autonomous-organization.md:1337	scholz 2016	Consolidated pre-audit reference used only the title <i>Platform Cooperativism</i> and omitted the publisher.	general reference	The official report, p. 14, and publication record identify <i>Platform Cooperativism: Challenging the Corporate Sharing Economy</i> and Rosa Luxemburg Stiftung—New York Office.	Complete the title and publisher while retaining the work as scoped context.

B043 — Ours to Hack and to Own

- **Canonical citation:** Scholz, T., & Schneider, N. (Eds.). (2016). *Ours to Hack and to Own: The Rise of Platform Cooperativism, a New Vision for the Future of Work and a Fairer Internet*. OR Books.
- **Metadata status:** verified; complete title and editor roles from OR Books; 2016 first-edition year from OCLC, while the current OR Books page displays January 2017
- **Primary record:** [OR Books publisher record](#); [OCLC first-edition catalog record](#)
- **Disposition:** retained
- **Defect tags:** ours-to-hack-subtitle
- **Forbidden live strings:** Scholz, T., & Schneider, N. (2016). *Ours to Hack and to Own: Platform cooperativism.*; Scholz & Schneider (2016). *Ours to Hack and to Own.*; Scholz, T., & Schneider, N. (2016). *Ours to Hack and to Own: The Rise of Platform Cooperativism, a New Vision for the Future of Work and a Fairer Internet*. OR Books.; Scholz, T., & Schneider, N. (2016). *Ours to Hack and to Own: Platform Cooperativism*. OR Books.; Scholz, T., & Schneider, N. (Eds.). (2016). *Ours to Hack and to Own: Platform Cooperativism*. OR Books.; Scholz, T., & Schneider, N. (Eds.). (2016). *Ours to Hack and to Own: The Rise of Platform Cooperativism*. OR Books.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B043-O001	live	docs/04-economics/00-overview.md:8	scholz schneider 2016	Platform-cooperative ownership debates form part of the governance premise.	supported	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B043-O002	live	docs/04-economics/01-trickle-out-investment-flow.md:88	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O003	live	docs/04-economics/02-revenue-allocation-framework.md:94	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O004	live	docs/13-member-trust-union/01-multi-level-trust-architecture.md:111	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O005	live	docs/13-member-trust-union/02-financial-product-design.md:161	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O006	live	docs/13-member-trust-union/03-governance-and-risk-distribution.md:153	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O007	live	docs/13-member-trust-union/05-member-experience-and-compensation.md:124	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O008	live	docs/13-member-trust-union/06-integration-with-icn-and-external-markets.md:141	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O009	live	docs/13-member-trust-union/08-mtu-as-domain-specific-hao-implementation.md:117	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O010	live	docs/humanized-autonomous-organization.md:5192	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O011	live	docs/humanized-autonomous-organization.md:5361	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O012	live	docs/humanized-autonomous-organization.md:5514	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O013	live	docs/humanized-autonomous-organization.md:5783	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B043-O014	live	docs/humanized-autonomous-organization.md:5934	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O015	live	docs/humanized-autonomous-organization.md:6204	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O016	live	docs/humanized-autonomous-organization.md:636	scholz schneider 2016	Platform-cooperative ownership debates form part of the governance premise.	supported	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
B043-O017	live	docs/humanized-autonomous-organization.md:736	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O018	live	docs/humanized-autonomous-organization.md:833	scholz schneider 2016	Reference-list placement provides bounded context on platform-cooperative ownership and governance.	general reference	OR Books publisher abstract describes cooperative ownership and governance of platforms.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B043-O019	removed	docs/04-economics/00-overview.md:8	scholz schneider 2016	The pre-audit sentence said the whole economic model draws on cooperative economics, attaching the edited volume to the ensuing HAO finance mechanics.	partially supported	The OR Books publisher abstract supports cooperative ownership and governance of platforms, but not the corpus's capital-allocation mechanics.	Limit the attribution to the platform-cooperative ownership premise and identify the economic mechanics as corpus proposals.
B043-O020	removed	docs/humanized-autonomous-organization.md:636	scholz schneider 2016	The consolidated pre-audit sentence said the economic model is guided by regenerative finance and cooperative economics.	partially supported	The OR Books publisher abstract supports cooperative ownership and governance of platforms, but not regenerative finance or the corpus's capital-allocation mechanics.	Limit the attribution to the platform-cooperative ownership premise and identify the economic mechanics as corpus proposals.
B043-O021	removed	docs/04-economics/01-trickle-out-investment-flow.md:88	scholz schneider 2016	The pre-audit reference used the noncanonical subtitle <i>Platform cooperativism</i> and omitted editor roles and publisher.	general reference	OR Books supplies the complete subtitle and editor roles; OCLC records the 2016 first edition.	Normalize editor roles, complete subtitle, and publisher while retaining scoped context.
B043-O022	removed	docs/humanized-autonomous-organization.md:736	scholz schneider 2016	Consolidated pre-audit duplicate used the noncanonical subtitle and omitted editor roles and publisher.	general reference	OR Books supplies the complete subtitle and editor roles; OCLC records the 2016 first edition.	Normalize editor roles, complete subtitle, and publisher while retaining scoped context.
B043-O023	removed	docs/04-economics/02-revenue-allocation-framework.md:94	scholz schneider 2016	The pre-audit reference omitted editor roles, subtitle, and publisher.	general reference	OR Books supplies the complete subtitle, editor roles, and publisher; OCLC records the 2016 first edition.	Normalize the complete canonical citation while retaining scoped context.
B043-O024	removed	docs/humanized-autonomous-organization.md:833	scholz schneider 2016	Consolidated pre-audit duplicate omitted editor roles, subtitle, and publisher.	general reference	OR Books supplies the complete subtitle, editor roles, and publisher; OCLC records the 2016 first edition.	Normalize the complete canonical citation while retaining scoped context.
B043-O025	removed	docs/13-member-trust-union/01-multi-level-trust-architecture.md:113	scholz schneider 2016	The pre-audit reference carried the complete subtitle but omitted the editors' role.	general reference	OR Books identifies Scholz and Schneider as editors; OCLC records the 2016 first edition.	Add the editor responsibility while retaining scoped context.
B043-O026	removed	docs/humanized-autonomous-organization.md:5194	scholz schneider 2016	Consolidated pre-audit duplicate carried the complete subtitle but omitted the editors' role.	general reference	OR Books identifies Scholz and Schneider as editors; OCLC records the 2016 first edition.	Add the editor responsibility while retaining scoped context.
B043-O027	removed	docs/13-member-trust-union/02-financial-product-design.md:161	scholz schneider 2016	The pre-audit reference used the noncanonical subtitle <i>Platform Cooperativism</i> and omitted editor roles.	general reference	OR Books supplies the complete subtitle and identifies Scholz and Schneider as editors.	Normalize editor roles and complete subtitle while retaining scoped context.
B043-O028	removed	docs/humanized-autonomous-organization.md:5367	scholz schneider 2016	Consolidated pre-audit duplicate used the noncanonical subtitle and omitted editor roles.	general reference	OR Books supplies the complete subtitle and identifies Scholz and Schneider as editors.	Normalize editor roles and complete subtitle while retaining scoped context.
B043-O029	removed	docs/13-member-trust-union/03-governance-and-risk-distribution.md:153	scholz schneider 2016	The pre-audit reference identified the editors but used the noncanonical subtitle <i>Platform Cooperativism</i> .	general reference	OR Books supplies the complete subtitle and editor roles.	Complete the subtitle while retaining scoped context.
B043-O030	removed	docs/humanized-autonomous-organization.md:5524	scholz schneider 2016	Consolidated pre-audit duplicate identified the editors but used the noncanonical subtitle.	general reference	OR Books supplies the complete subtitle and editor roles.	Complete the subtitle while retaining scoped context.
B043-O031	removed	docs/13-member-trust-union/05-member-experience-and-compensation.md:124	scholz schneider 2016	The pre-audit reference truncated the subtitle after <i>The Rise of Platform Cooperativism</i> .	general reference	OR Books supplies the complete subtitle ending a <i>New Vision for the Future of Work and a Fairer Internet</i> .	Complete the subtitle while retaining scoped context.
B043-O032	removed	docs/13-member-trust-union/06-integration-with-icn-and-external-markets.md:141	scholz schneider 2016	The pre-audit reference truncated the subtitle after <i>The Rise of Platform Cooperativism</i> .	general reference	OR Books supplies the complete subtitle ending a <i>New Vision for the Future of Work and a Fairer Internet</i> .	Complete the subtitle while retaining scoped context.
B043-O033	removed	docs/13-member-trust-union/08-mtu-as-domain-specific-hao-implementation.md:117	scholz schneider 2016	The pre-audit reference truncated the subtitle after <i>The Rise of Platform Cooperativism</i> .	general reference	OR Books supplies the complete subtitle ending a <i>New Vision for the Future of Work and a Fairer Internet</i> .	Complete the subtitle while retaining scoped context.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B043-0034	removed	docs/humanized-autonomous-organization.md:5795	scholz schneider 2016	Consolidated pre-audit reference truncated the subtitle after <i>The Rise of Platform Cooperativism</i> .	general reference	OR Books supplies the complete subtitle ending a <i>New Vision for the Future of Work and a Fairer Internet</i> .	Complete the subtitle while retaining scoped context.
B043-0035	removed	docs/humanized-autonomous-organization.md:5946	scholz schneider 2016	Consolidated pre-audit reference truncated the subtitle after <i>The Rise of Platform Cooperativism</i> .	general reference	OR Books supplies the complete subtitle ending a <i>New Vision for the Future of Work and a Fairer Internet</i> .	Complete the subtitle while retaining scoped context.
B043-0036	removed	docs/humanized-autonomous-organization.md:6219	scholz schneider 2016	Consolidated pre-audit reference truncated the subtitle after <i>The Rise of Platform Cooperativism</i> .	general reference	OR Books supplies the complete subtitle ending a <i>New Vision for the Future of Work and a Fairer Internet</i> .	Complete the subtitle while retaining scoped context.

B044 — Seeing Like a State

- **Canonical citation:** Scott, J. C. (1998). *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press.
- **Metadata status:** verified against the linked primary record
- **Primary record:** [Yale official publication record](#)
- **Disposition:** retained
- **Defect tags:** none
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B044-0001	live	docs/13-member-trust-union/01-multi-level-trust-architecture.md:109	scott 1998	Reference-list placement provides bounded context on the limits of central planning and loss of local knowledge.	general reference	Yale publisher abstract cautions that central schemes can erase local knowledge and interdependencies.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B044-0002	live	docs/humanized-autonomous-organization.md:5190	scott 1998	Reference-list placement provides bounded context on the limits of central planning and loss of local knowledge.	general reference	Yale publisher abstract cautions that central schemes can erase local knowledge and interdependencies.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

B045 — Touchpoint, volume 11, number 1

- **Canonical citation:** Service Design Network. (2019). *Touchpoint: The Journal of Service Design*, 11(1), Service Design for Innovation and Start-ups.
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** [Service Design Network issue record](#)
- **Disposition:** retained
- **Defect tags:** touchpoint-year-scope
- **Forbidden live strings:** Service Design Network. (2020) ; (Service Design Network, 2020) ; Service Design Network (2020)

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B045-0001	live	docs/13-member-trust-union/04-technology-and-physical-digital-integration.md:135	service design network 2019	Reference-list placement provides bounded context on service-design and innovation practice.	general reference	Official issue abstract establishes service-design and innovation scope, not hospitality-based community banking.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B045-0002	live	docs/humanized-autonomous-organization.md:5655	service design network 2019	Reference-list placement provides bounded context on service-design and innovation practice.	general reference	Official issue abstract establishes service-design and innovation scope, not hospitality-based community banking.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B045-0003	removed	docs/13-member-trust-union/01-multi-level-trust-architecture.md:109	service design network 2020	Historical pre-audit occurrence using the superseded or removed service design network 2020 attribution.	unsupported	The official issue abstract and contents do not establish hospitality-driven community-bank architecture.	Remove this attribution and its section bibliography anchor.
B045-0004	removed	docs/13-member-trust-union/01-multi-level-trust-architecture.md:44	service design network 2020	Historical pre-audit occurrence using the superseded or removed service design network 2020 attribution.	unsupported	The official issue abstract and contents do not establish hospitality-driven community-bank architecture.	Remove this attribution and its section bibliography anchor.
B045-0005	removed	docs/13-member-trust-union/04-technology-and-physical-digital-integration.md:135	service design network 2020	Historical pre-audit occurrence using the superseded or removed service design network 2020 attribution.	general reference	The official issue record supports general service-design and innovation context and dates volume 11, number 1 to July 2019; it does not support hospitality-driven community banking.	Supersede with the corrected 2019 identity only in the technology section.
B045-0006	removed	docs/humanized-autonomous-organization.md:5119	service design network 2020	Historical pre-audit occurrence using the superseded or removed service design network 2020 attribution.	unsupported	The official issue abstract and contents do not establish hospitality-driven community-bank architecture.	Remove this attribution and its section bibliography anchor.

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B045-O007	removed	docs/humanized-autonomous-organization.md:5190	service design network 2020	Historical pre-audit occurrence using the superseded or removed service design network 2020 attribution.	unsupported	The official issue abstract and contents do not establish hospitality-driven community-bank architecture.	Remove this attribution and its section bibliography anchor.
B045-O008	removed	docs/humanized-autonomous-organization.md:5667	service design network 2020	Historical pre-audit occurrence using the superseded or removed service design network 2020 attribution.	general reference	The official issue record supports general service-design and innovation context and dates volume 11, number 1 to July 2019; it does not support hospitality-driven community banking.	Supersede with the corrected 2019 identity only in the technology section.

B046 — The Precariat

- **Canonical citation:** Standing, G. (2011). *The Precariat: The New Dangerous Class*. Bloomsbury Academic. <https://doi.org/10.5040/9781849664554>
- **Metadata status:** verified against the linked primary record
- **Primary record:** [Bloomsbury Collections 2011 record](#)
- **Disposition:** retained
- **Defect tags:** none
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B046-O001	live	docs/13-member-trust-union/05-member-experience-and-compensation.md:130	standing 2011	Reference-list placement provides bounded context on labor insecurity and precarious work.	general reference	Original chapter 1 describes insecure labor relations and absent occupational identity.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B046-O002	live	docs/humanized-autonomous-organization.md:5789	standing 2011	Reference-list placement provides bounded context on labor insecurity and precarious work.	general reference	Original chapter 1 describes insecure labor relations and absent occupational identity.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

B047 — Trust as Infrastructure

- **Canonical citation:** No canonical Thomas and Mainwaring (2018) SSIR publication was established for the baseline string.
- **Metadata status:** not currently verifiable; the supplied SSIR URL returns 404 and the bounded first-party, registry, and archive search located no official title-and-byline record
- **Primary record:** No primary record was located. The [supplied SSIR path](#) returns 404; SSIR's own search/index and the documented title, byline, author, DOI/Crossref, and archive checks yielded no official record or primary text.
- **Disposition:** removed
- **Defect tags:** thomas-mainwaring-byline
- **Forbidden live strings:** Thomas, R., & Mainwaring, S. (2018). *Trust as infrastructure: Redesigning financial systems for community wealth*. *Stanford Social Innovation Review*. ; https://ssir.org/articles/entry/trust_as_infrastructure

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B047-O001	removed	docs/13-member-trust-union/01-multi-level-trust-architecture.md:117	thomas mainwaring 2018	Pre-audit reference to the uncorroborated Thomas and Mainwaring SSIR identity.	not verifiable	The supplied SSIR path returns 404, and the search could not locate an official title/byline record or primary text to support the trust-architecture section.	Remove the placement without transferring it to a similar source. Restore only if an archived first-party identity and claim source are obtained.
B047-O002	removed	docs/13-member-trust-union/02-financial-product-design.md:163	thomas mainwaring 2018	Pre-audit reference to the uncorroborated Thomas and Mainwaring SSIR identity.	not verifiable	The supplied SSIR path returns 404, and the search could not locate an official title/byline record or primary text to support the financial-product section.	Remove the placement without transferring it to a similar source. Restore only if an archived first-party identity and claim source are obtained.
B047-O003	removed	docs/humanized-autonomous-organization.md:5198	thomas mainwaring 2018	Consolidated duplicate of the uncorroborated Thomas and Mainwaring SSIR identity.	not verifiable	The supplied SSIR path returns 404, and the search could not locate an official title/byline record or primary text to support the trust-architecture section.	Remove the placement without transferring it to a similar source. Restore only if an archived first-party identity and claim source are obtained.
B047-O004	removed	docs/humanized-autonomous-organization.md:5369	thomas mainwaring 2018	Consolidated duplicate of the uncorroborated Thomas and Mainwaring SSIR identity.	not verifiable	The supplied SSIR path returns 404, and the search could not locate an official title/byline record or primary text to support the financial-product section.	Remove the placement without transferring it to a similar source. Restore only if an archived first-party identity and claim source are obtained.

B048 — The Problem with Work

- **Canonical citation:** Weeks, K. (2011). *The Problem with Work: Feminism, Marxism, Antiwork Politics, and Postwork Imaginaries*. Duke University Press. <https://doi.org/10.1215/9780822394723>
- **Metadata status:** verified against the linked primary record
- **Primary record:** [Duke University Press book record](#)
- **Disposition:** retained
- **Defect tags:** none
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
B048-O001	live	docs/13-member-trust-union/05-member-experience-and-compensation.md:132	weeks 2011	Reference-list placement provides bounded context on critiques of wage work and postwork imaginaries.	general reference	Duke publisher abstract challenges wage work as the primary basis of income and identity.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
B048-O002	live	docs/humanized-autonomous-organization.md:5791	weeks 2011	Reference-list placement provides bounded context on critiques of wage work and postwork imaginaries.	general reference	Duke publisher abstract challenges wage work as the primary basis of income and identity.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.

S001 — Neocortex size as a constraint on group size in primates

- **Canonical citation:** Dunbar, R. I. M. (1992). Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*, 22(6), 469–493. [https://doi.org/10.1016/0047-2484\(92\)90081-J](https://doi.org/10.1016/0047-2484(92)90081-J)
- **Metadata status:** verified; current corpus placement removed after claim review
- **Primary record:** [Elsevier article record](#)
- **Disposition:** removed
- **Defect tags:** none
- **Forbidden live strings:** consistent with Dunbar's theory of cognitive limits on stable social relationships [Dunbar, 1992]; This reflects Dunbar's theory of cognitive limits on stable social relationships [Dunbar, 1992]

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
S001-O001	removed	docs/12-icn/01-structural-features.md:16	dunbar 1992	The pre-audit chapter said an 8–15-contributor range was consistent with Dunbar's theory and intended to support trust, agility, and internal coherence.	unsupported	The inspected article, pp. 469–493, does not test or recommend an 8–15-person human work team.	Remove the Dunbar attribution; retain the range only as an explicit corpus heuristic.
S001-O002	removed	docs/humanized-autonomous-organization.md:4333	dunbar 1992	The consolidated pre-audit text said the 8–15-contributor range maximizes trust, agility, and coherence and reflects Dunbar's theory.	unsupported	The inspected article, pp. 469–493, does not test or recommend an 8–15-person human work team.	Remove the Dunbar attribution; retain the range only as an explicit corpus heuristic.

S002 — Origins and Development of the Cellular Organization

- **Canonical citation:** Schley, D. G. (2010). Origins and development of the cellular organization. In M. A. Rahim (Ed.), *Current Topics in Management: Volume 14, Organizational Behavior, Performance, and Effectiveness* (pp. 119–148). Transaction Publishers. <https://doi.org/10.4324/9780203793985-6>
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** [Taylor & Francis chapter record](#)
- **Disposition:** retained
- **Defect tags:** schley-year
- **Forbidden live strings:** Schley (2009)

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
S002-O001	live	docs/appendices/lineage.md:74	schley 2010	The lineage footnote is a contextual annotation to work on cellular organization, not evidenced ancestry.	supported	Publisher chapter abstract and DOI establish a 2010 chapter on cellular organization; commit 7863f9ac adds text and footnote together.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
S002-O002	live	docs/appendices/lineage.md:78	schley 2010	Reference-list placement provides bounded context on cellular organization.	general reference	Publisher chapter abstract and DOI establish a 2010 chapter on cellular organization; commit 7863f9ac adds text and footnote together.	Retain as scoped context; do not treat it as support for the section's HAO-specific mechanisms.
S002-O003	removed	docs/appendices/lineage.md:72	schley 2009	The pre-audit lineage said the Schley footnote was added after the sibling-site text and was not drawn on to write it.	partially supported	The publisher chapter abstract and DOI establish a 2010 chapter; sibling-site commit 7863f9ac adds the definition and footnote together, contradicting the recorded “added afterward” chronology.	Correct the year and state that the annotation entered with the text in available Git history; do not promote it to evidenced ancestry.

S003 — Managing Leading-Edge Multinational Corporations

- **Canonical citation:** Snow, C. C., Miles, R. E., & Allred, B. B. (2003). Managing leading-edge multinational corporations. In B. McKern (Ed.), *Managing the Global Network Corporation* (digital ed., pp. 215–233). Routledge. <https://doi.org/10.4324/9780203478813-21>
- **Metadata status:** verified; pre-audit metadata repaired
- **Primary record:** [Taylor & Francis chapter record](#)
- **Disposition:** retained
- **Defect tags:** snow-author-order
- **Forbidden live strings:** Allred, Miles, and Snow (2003)

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
S003-O001	live	docs/appendices/lineage.md:74	snow miles allred 2003	The lineage footnote is a contextual annotation to work on cellular organizations, not evidenced ancestry.	supported	Publisher chapter abstract and DOI establish the Snow–Miles–Allred chapter; commit 7863f9ac adds text and footnote together.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
S003-O002	live	docs/appendices/lineage.md:79	snow miles allred 2003	Reference-list placement provides bounded context on cellular and networked organizational forms.	general reference	Publisher chapter abstract and DOI establish the Snow–Miles–Allred chapter; commit 7863f9ac adds text and footnote together.	Retain as scoped context; do not treat it as support for the section’s HAO-specific mechanisms.
S003-O003	removed	docs/appendices/lineage.md:72	allred miles snow 2003	The pre-audit lineage said the Allred, Miles, and Snow footnote was added after the sibling-site text and was not drawn on to write it.	partially supported	The publisher chapter abstract and DOI establish the Snow–Miles–Allred chapter; sibling-site commit 7863f9ac adds the definition and footnote together, contradicting the recorded “added afterward” chronology.	Correct the author order and state that the annotation entered with the text in available Git history; do not promote it to evidenced ancestry.

S004 — Cellular organizational structure, permanent revision 756005860

- **Canonical citation:** Wikipedia contributors. (2016). Cellular organizational structure. *Wikipedia*, permanent revision 756005860, 21 December 2016, 12:31:29 UTC. Cited section: “Background and concepts.”
- **Metadata status:** verified against the linked primary record
- **Primary record:** [Permanent revision 756005860](#)
- **Disposition:** retained
- **Defect tags:** none
- **Forbidden live strings:** none

Occurrence	State	Location	Fingerprint	Claim	Verdict	Evidence	Action
S004-O001	live	docs/appendices/lineage.md:72	wikipedia contributors 2016	The local artifacts reproduce wording traceable to Wikipedia revision 756005860; tracked chronology has stated limits.	supported	Wikipedia section Background and concepts, revision 756005860, and the cited Git commits establish the bounded chronology.	Retain the repaired, bounded attribution; the HAO-specific mechanism remains corpus-authored.
S004-O002	live	docs/appendices/lineage.md:80	wikipedia contributors 2016	Reference-list placement provides bounded context on the exact wording reproduced by the local source artifacts.	general reference	Wikipedia section Background and concepts, revision 756005860, and the cited Git commits establish the bounded chronology.	Retain as scoped context; do not treat it as support for the section’s HAO-specific mechanisms.
S004-O003	removed	docs/appendices/lineage.md:72	wikipedia	The pre-audit lineage correctly traced the passage to Wikipedia but said it was later copied into the sibling site and that its footnotes were added afterward.	partially supported	Wikipedia section Background and concepts, revision 756005860, establishes the wording; sibling-site commit 7863f9ac adds text and footnotes together, while repository dates cannot establish off-repository drafting order.	Retain the Wikipedia origin, remove the contradicted footnote chronology, and state the remaining direction-of-copying limit.

HAO Lineage and Vocabulary Genealogy

A record of where the model came from, which of its earlier terms became current ones, and which were dropped.

1. Purpose

This appendix records the documented history of the Humanized Autonomous Organization (HAO) — the network’s coordinating framework — as an organizational-design model: the account its author gave of how the work began, the bodies of prior work it draws on, and the vocabulary it used before the current one.

The current vocabulary of the corpus — HAO, United Micro Enterprise (UME) — a small, self-managing venture team ($\leq \sim 15$ people), Strategic Enterprise Partnership (SEP) — a joint venture between teams, and Micro Enterprise Ecosystem (MEE) — the network’s protected internal economy — replaced an earlier vocabulary built on cellular metaphors. The earlier vocabulary is recorded here, in §4 and §5, as history. It is not live vocabulary, and the working chapters do not use it. `docs/appendices/terminology.md` remains the corpus’s authoritative glossary; where this appendix and that glossary differ, the glossary governs current usage and this appendix records what came before it.

Sources and method. The initial draft drew on a staged extraction over material outside the corpus: a dated forum post by the author, a 19-entry lexicon on a sibling documentation site, two lexicon entries drafted in an AI interview transcript but never promoted to standalone files, a research index page, and a set of AI interview transcripts. Later targeted checks reopened primary material, including the prompt-and-response pair discussed in §2.2 and the three Beyond Budgeting works compared in §3.4. Every claimed successor term in §4 was separately checked against the live corpus in `~/code/hao/docs/`, and three points where the staged extraction no longer matches the live corpus are marked where they occur.

On quotation. The corpus does not quote individuals. This appendix is the single exception: the author’s own recorded words are quoted here, sparingly, and each quotation is marked as the author’s. The value of the origin account is that it is a first-person record, so it is not paraphrased into the corpus register. Everything outside the marked quotations is this appendix’s own neutral-register summary. Material generated by AI systems in the source transcripts is summarized rather than quoted, since it is neither the author’s words nor a citable publication. Quotations follow the primary sources exactly, including their terminal punctuation; where a quotation contains a quotation of its own, the inner marks are set as single quotes for nesting, and no other change is made to any quoted text. Where a quotation stops short of the end of its passage, or resumes on a later line of the source, the appendix says so rather than joining the parts silently.

2. The author’s origin account

Source: `~/code/site-provide-coop/provide.coop/bits/tims-co-breakdown.txt`, dated 2024-04-28, a forum or chat post. The sequence below follows the post as written. All quotations in this section are the corpus author’s own words from that post.

2.1 Motive and the earlier coinage

The author introduces himself as a staff DevOps engineer at a firm owned by private equity, previously employed at a non-profit whose board decided, in his words, to “sell all IP and assets and pivot focus” before arranging for staff to be hired by the acquirer. He gives that as the occasion for the post.

The idea is described as a centralized and decentralized network of intrinsically linked entities, engineered to implement polycentric governance principles so that the ecosystem acts as a mesh network for its participants. He then names it. In the author’s words: “I’m calling it a Cellular Cooperative.”

This is the earliest recorded name for the model. It is not the name the corpus uses.

The author states the multi-year motivation behind the project in a single sentence, in his own words: “Make owning a business as easy as applying for a job. For a wide range of industries.” Everything else in the post is offered as one part of that goal.

2.2 Encountering the term “polycentric governance”

The post records how the term entered the vocabulary. In the author’s words: “FWIW, I had no idea what ‘polycentric governance’ was until like a month ago... when I managed to write the right prompt to get one of those LLM AIs to inform me that PG is what I’ve been thinking about but didn’t have the words to discuss.”

The term is load-bearing throughout the rest of the corpus. The author’s stated route of entry is a language model exchange roughly a month before the post, rather than direct engagement with the polycentricity literature. A surviving local pair in `~/code/papr/polycentric-governance/` consists of `001-describe-polycentric.md`, a request to describe polycentric governance, and `001-r-chatgpt.md`, a response that names Elinor and Vincent Ostrom. Both files carry local modification dates of 20 March 2024 and first enter the `papr` Git record in commit `52bd264` on 21 March 2024. The dates establish when the preserved files were recorded, not the time of the original conversation. Because the prompt already uses the term, the pair does not establish when the author first encountered it. Nor does the response establish that he read the Ostroms’ publications. §3.1 records how this relates to the retrospective attribution of the concept to Ostrom.

2.3 The unit and the sociocratic circle

The post defines three units in sequence.

- **Cellular Organization (CO):** a wrapper that coordinates Cells, which the author states may itself be a publicly traded company, a non-profit, a cooperative, a school, or an institution. He distinguishes it from two named operating models. In the author’s words: “A CO is *not* Holacracy, or Sociocracy. But, both of those Operating Models can be implemented within a CO.”
- **Cell / Cellular Unit (CU):** an intentionally small group of actors, resources, stakeholders, and shared mission that produces and provides value. The relationship to sociocracy is stated directly. In the author’s words: “A ‘Cell’ is akin to a ‘Circle’ in Sociocracy. But more abstract.”
- **Cellular Venture (CV):** a group of Cells that have come together to produce shared value, which the author states implies that the Cells have defined and agreed on an Operating Agreement. Cells and actors inside a CV choose which governance model applies to them, from sociocracy, holacracy, or hierarchy.

A mandated cross-functional team is dismissed as a degenerate case of the CV: the author describes it as a hierarchy inside a hierarchy, for which he coins the acronym HIAH, and states that no autonomy exists within that structure.

The composition claim closes the definitional sequence. In the author’s words: “People -> Cell -> Cellular Venture -> Cellular Organization. Nothing really novel in that. It’s Sociocracy. :)”

A further group of terms is named in the post but left undefined, listed as more operational than structural: Cellular Integration, Cellular Lifecycle Management, Cellular Culture, Cellular Collapse, Cellular Degradation, Dynamic Contracts, and Dynamic Operating Agreements. Each is traced in §4.

2.4 The choice of limited-liability form over cooperative entity form

The post rejects the cooperative legal form explicitly. In the author’s words: “Also, this isn’t designed to take advantage of traditional Cooperative entity structures. It’s meant to be implemented in a Limited Liability Company, or a Series LLC.” The reason follows on the next line of the post, again in the author’s words: “That way each Cell has its own corporate veil.”

The stated design goal is a liability boundary per unit. Layering limited-liability entities, or using a series structure, gives each unit its own corporate veil in a way that a single cooperative entity holding all activity does not. The cooperative framing is retained at the level of culture and governance norms — worker ownership and democratic decision-making are stated as Cell characteristics on the sibling site’s `cell1.mdx` — while the legal wrapper is deliberately something else. §3.6 records the same split for the cooperative movement as an influence.

2.5 Stated scope and stated limits

Two limits on the claim are recorded in the post itself.

The author anticipates readers proposing technology solutions and excludes them. In the author’s words: “So, no blockchain or DLT or web3 or any of that stuff is within the scope of what I’m sharing.” He characterizes what he is sharing as the architecture of a blueprint representing only the exterior of a home.

He also disclaims originality and authority for the vocabulary. In the author’s words: “Please keep in mind, there is no ‘official’ definition for these terms that I’ve been able to find, within this organizational metaphor.” The passage resumes on the next line of the post, again in the author’s words: “So... I’m kinda making up the terms as I go along. And I know, some of the concepts are likely labeled by some other word, or damn similar. :)”

That disclaimer bears on the rest of this appendix. The author does not claim the vocabulary was novel, only that he had not located prior art for it at the time of writing. The retrospective identifications of prior art recorded in §3 were made later.

2.6 Stated research method

The post describes the author’s validation approach: describing problems to language models and treating their predictions of his intended solutions as confirmation of his reasoning. It then reproduces two responses from a language model analyzing the Cellular Cooperative concept against cooperatives, holacracy and sociocracy, and agile teams, naming gaps in scaling democratic process, coordination overhead, and incentive design, and naming claimed innovations in hybrid governance, collective intelligence, and resilience.

On this appendix’s reading, the transcript is the earliest recorded outside review of the model; the author records the method but does not characterize the exchange this way. As a single language model’s response to the author’s own framing, it is not independent verification, and this appendix does not treat it as evidence for any claim beyond the fact that the exchange occurred.

The post and the sibling site’s `cellular-organization.mdx` reproduce wording traceable to the `Background and concepts` section of Wikipedia’s permanent revision `756005860` (Wikipedia contributors, 2016). The sibling site’s available Git history records the definition and its two scholarly footnotes entering together in commit `7863f9ac` on 11 April 2024; the archived-post file enters that repository in commit `14e2c882` on 28 April 2024. Those repository dates do not establish when either artifact was drafted outside the repository or which local artifact supplied the other.

The two footnotes are contextual annotations to work on cellular organization (Schley, 2010; Snow, Miles, & Allred, 2003). They were not added after the sibling-site definition in the recorded file history. Neither the annotations nor that chronology establishes that either scholarly work influenced development of the HAO model, so §3 does not promote them to evidenced ancestry.

References for this chronology

- Schley, D. G. (2010). Origins and development of the cellular organization. In M. A. Rahim (Ed.), *Current Topics in Management: Volume 14, Organizational Behavior, Performance, and Effectiveness* (pp. 119–148). Transaction Publishers. <https://doi.org/10.4324/9780203793985-6>
- Snow, C. C., Miles, R. E., & Allred, B. B. (2003). Managing leading-edge multinational corporations. In B. McKern (Ed.), *Managing the Global Network Corporation* (digital ed., pp. 215–233). Routledge. <https://doi.org/10.4324/9780203478813-21>
- Wikipedia contributors. (2016). Cellular organizational structure. *Wikipedia*, permanent revision 756005860, 21 December 2016, 12:31:29 UTC. Cited section: “Background and concepts.” https://en.wikipedia.org/w/index.php?title=Cellular_organizational_structure&oldid=756005860

3. Intellectual ancestry

Each influence below is marked **evidenced** or **asserted**.

- **Evidenced** means the connection is shown in the primary sources reviewed for the staged extraction: the source states what was taken and where it came from, at or near the point where the borrowed idea is used.

- **Asserted** means the connection is claimed — most often on the research index page at `~/code/site-provide-coop/provide.coop/research/index.mdx` — without support traced to the author’s working notes, and without the borrowed content appearing anywhere the claim could be checked against.

Two of the six are evidenced; four are asserted. The four asserted influences share a pattern: each is named on a research index page written after the model existed, identifying prior art for work already done, rather than a source the model was built from.

3.1 Ostrom and polycentricity — asserted

What is claimed to have been taken: the governing vocabulary and the conceptual root of polycentric governance as the decision-rights model for the Integrated Cooperative Network (ICN) — the reference cooperative business network.

Status: asserted. The research index page states that Ostrom-lineage work on multiple centers of authority is where the ICN’s governance model originates. That page states the derivation; it does not argue it, cite a specific passage, or trace a mechanism from the literature into the model.

The author’s own origin account gives a different route for the term. As recorded in §2.2, he describes discovering the phrase “polycentric governance” through a language model prompt about a month before writing, not through reading Ostrom. The two accounts are not strictly contradictory: a term can be encountered indirectly and its literature read afterward. But the research page’s framing, that this is where the model originates, describes a scholarly grounding identified after the fact rather than the route by which the vocabulary was actually acquired, and no source reviewed records the reading that would close the gap.

The term survived the vocabulary change intact. “Polycentric” remains load-bearing throughout `docs/`, and the Adaptive Governance Framework (AGF) — the network’s layered governance system — is defined in `docs/appendices/terminology.md` §II as a multi-level, polycentric governance model.

3.2 Beer’s Viable System Model and cybernetics — asserted

What is claimed to have been taken: a systems-theory basis for cellular organization generally.

Status: asserted. The research index page names Beer’s Viable System Model and the systems-theory traditions underneath cellular organization. A separate passing link to the Wikipedia article on variety in cybernetics appears in `~/code/provide.io/content/english/governance.md`, offered as an aside rather than as a derivation.

No content specific to the Viable System Model — recursion, the viability criterion, or the five named systems — appears in the origin account, in any lexicon entry, or on the governance page. The connection is named on one index page and developed nowhere. It is the weakest-supported of the six.

3.3 Holonic manufacturing — asserted, and named in the sources as the closest prior art

What is claimed to have been taken: nothing is claimed to have been taken. Holonic manufacturing is named in the sources as the closest prior art to the model’s basic unit, the Cell, rather than as a source the Cell was derived from.

Status: asserted, and retrospective. The research index page describes an existing engineering literature on autonomous units that cooperate without central control, and observes that it comes from factories rather than from cooperatives. The page states this identification; it does not argue it. No comparison of the holon to the Cell is set out, no property of one is mapped onto the other, and no source from that literature is cited at the point of the claim.

The identification is retrospective on the evidence available. Holonic manufacturing is not mentioned in `tims-co-breakdown.txt` or in any lexicon entry, and the author does not cite it at the point where he coins the Cell. It appears only on the later research index page. Read together with the author’s own disclaimer in §2.5 — that he expected the concepts were already labeled by other words — the identification is best read as a later answer to a question the author had already posed about his own vocabulary, not as an unacknowledged debt.

3.4 Beyond Budgeting — asserted

What is claimed to have been taken: devolved financial authority, framed on the research index page as the question of how an autonomous unit obtains money without asking permission.

Status: asserted, and retrospective. The claim rests on the research index page alone and is stated rather than argued. Nothing corresponding to it appears in the origin materials. The status matches holonic manufacturing.

Beyond Budgeting is the one asserted influence with substantial representation in the assembled research library: `docs/appendices/references.md` §4 lists five works on it. Those works were collected; no source reviewed shows any of them being used.

Research comparison, updated 2026-09-24. Three directly rechecked works describe devolved decision-making, flexible allocation and relative targets among Beyond Budgeting adopters, as well as disagreement over whether the approach must be adopted as a complete model. Hudson’s study reports partial adoption in practice. The other two collected works remain outside this recheck. This is a comparison with the HAO’s proposed arrangements, not evidence that the literature influenced their design or that the arrangements will work in this setting. The ancestry status remains asserted.

3.5 Sociocracy — evidenced

What was taken: the analogy between the basic unit and the sociocratic circle; the layering of Cell, Cellular Venture, and Cellular Organization; and the position that a Cellular Organization can host sociocracy or holacracy as an internal operating model without being either.

Status: evidenced. This is the most directly attested influence of the six, cited by the author in his own words at the moment of coining the terms, three times in a single post — the Cell-to-Circle analogy, the statement that a CO is neither holacracy nor sociocracy but can host both, and the composition claim that the four-level stack is sociocracy. All three are quoted in §2.3.

The corpus retains the mechanism under a different label. `docs/appendices/terminology.md` §II defines Consent Governance as a decision-making model focused on resolving objections rather than achieving majority rule, without naming sociocracy as its source.

3.6 The cooperative movement — evidenced at the level of culture, rejected at the level of legal form

What was taken: worker-ownership and democratic-decision norms at the level of unit culture, and the word “Cooperative” in the model’s original name.

What was declined: the cooperative entity form itself.

Status: evidenced, with the split recorded in the sources. Both halves are attested. The sibling site’s `cell.mdx` states that Cells are typically structured as worker-owned cooperatives, and the research index page catalogs a work on humanistic governance in democratic organizations as a source. Against that, `governance.md` and `tims-co-breakdown.txt` both state that the model is deliberately not built on cooperative legal entity structures, in favor of an LLC or Series LLC so that each unit carries its own corporate veil (§2.4).

The cooperative movement supplied the internal governance ethic. It did not supply the entity form. This is the only influence in this section where the sources record both what was adopted and what was declined, and the reason for declining it.

4. Vocabulary genealogy

The earlier vocabulary is drawn from the 19 files in `~/code/site-provide-coop/provide.coop/docs/lexicon/`, plus two entries drafted in AI interview transcripts and never promoted to standalone lexicon files (Mycorium and Holographic Organization), for 21 entries. Current terms were checked against `docs/appendices/terminology.md` and against targeted searches of `~/code/hao/docs/`.

Scope of the rename. The rename is specific to this corpus. `~/code/site-provide-coop/provide.coop/docs/icn/` is a separate documentation tree that remains live and continues to use “Cell” and “Cellular Venture” as of its most recent edits, dated 2026-08-04. A term marked below as renamed within this corpus may still be the current term on that site.

Status vocabulary used in the table. *Renamed* means both terms are defined and the definitions correspond. *Unconfirmed* means a successor is plausible on the definitions but no source states the equivalence. *Dropped, no successor* means no current term occupies the same role.

Earlier term	Current corpus term	Note
Cell / Cellular Unit (CU)	UME (United Micro Enterprise)	Renamed. Definitions correspond closely: a semi-autonomous, self-governing, small-scale unit. Verified in <code>terminology.md</code> §I. “Cell” is not retired network-wide; it remains the live term on the sibling ICN site.
Cellular Organization (CO)	HAO, as closest analogue only — unconfirmed	No rename is documented in any source. HAO fills a structurally similar role, a generic coordinating wrapper that can house different internal governance models, and is defined in <code>terminology.md</code> §I. No source states the equivalence. Treat as an unevidenced analogy.
Cellular Venture (CV)	SEP (Strategic Enterprise Partnership)	Renamed. <code>cellular-venture.mdx</code> defines a temporary collaboration of multiple Cells focused on a defined goal; <code>terminology.md</code> §I defines a SEP as a time-bound or goal-specific alliance between two or more UMEs. “Cellular Venture” also persists unchanged on the sibling site.
Cellular Culture	Dropped, no successor	No equivalent named entry exists in <code>terminology.md</code> . The ground is partly covered, under no unifying term, by <code>docs/07-human-systems/02-cultural-onboarding-and-ritual-design.md</code> .
Cellular Degradation	Enterprise Degradation	Renamed and formalized, not dropped. <code>terminology.md</code> §VI defines it as drift from HAO-aligned principles due to cultural, financial, or operational misalignment. <code>docs/06-lifecycle/03-collapse-and-containment-protocols.md</code> defines it as a pre-collapse state characterized by drift from core principles, reduced functionality, or social dysfunction. Both were checked directly; see §5.1.
Cellular Collapse	Enterprise Collapse, formalized as the Collapse and Containment Protocols	Renamed and substantially expanded, not dropped. <code>terminology.md</code> §VI defines a terminal state where a UME ceases operation. <code>docs/06-lifecycle/03-collapse-and-containment-protocols.md</code> adds five trigger classes, a seven-phase lifecycle from Detection to Reintegration, three containment tiers (CPT-1 to CPT-3), and post-collapse learning capture. See §5.1.
Cellular Integration	EIA (Enterprise Integration Assessment), as closest analogue — unconfirmed	<code>terminology.md</code> §II defines the EIA as a formalized process for evaluating whether and how external entities can be integrated into the ICN or HAO. <code>cellular-integration.mdx</code> describes a measured evaluation of how external entities can be divided into Cells and formed into linked Cellular Ventures. Conceptually close; no source cross-references the two. Treat as a likely but unconfirmed continuation.
Cellular Lifecycle Management (CLM)	UME Lifecycle Management	Label renamed, content changed. <code>terminology.md</code> §VI carries the renamed label. Chapter 06 later recorded the three-type framework as an exploratory layer over the uniform UME lifecycle; see §5.5.
Dynamic Contract (generic)	Absorbed into the DEA (Dynamic Enterprise Agreement) — a versioned operating agreement replacing fixed bylaws	The earlier entry was thin and self-qualifying, describing contracts that are not smart contracts but may be implemented as such, primarily concerning Operating Agreements and internal orderings. It functions as a precursor to the Dynamic Operating Agreement entry below rather than as a distinct survivor.
Dynamic Cross-Network Contracts	SEP Revenue Contract, as partial analogue — unconfirmed	<code>terminology.md</code> §III defines a dynamic revenue-sharing agreement specific to a SEP covering contribution, risk, trust weight, and revenue distribution. The earlier entry’s distinctive framing — safeguards against securities-law violations, designed to fail the Howey test — has no counterpart: a search of <code>docs/</code> for “Howey” returns nothing. Public-market exposure is addressed by a different route in <code>docs/08-legal/04-public-market-interaction-framework.md</code> .
Dynamic Operating Agreement	DEA (Dynamic Enterprise Agreement)	Renamed, directly. <code>governance.md</code> describes a versioned Operating Agreement deployed in place of an amendment process; <code>terminology.md</code> §II defines the DEA as a versioned, evolving governance charter that replaces fixed bylaws.
Polycentric Governance	Polycentric Governance, with the AGF (Adaptive Governance Framework) built on it	Survived. “Polycentric” remains in use throughout <code>docs/</code> , including in the HAO and AGF definitions in <code>terminology.md</code> §I and §II. Continued use plus a new named framework, rather than a rename.

Earlier term	Current corpus term	Note
Polycentric+ (“Polycentric Plus”)	Dropped, no successor	No such entry exists anywhere in <code>docs/</code> (searched, zero matches). Its content — proactive conflict-resolution tooling and supportive infrastructure — is diffusely present across <code>docs/07-human-systems/</code> under no unifying label. The earlier entry is self-deprecating about its own originality.
Polycentricity	Polycentricity	Survived unchanged. A generic term, used throughout.
Arm’s Length	Dropped as a named entry, no successor	The earlier entry reproduced a legal-dictionary definition verbatim and was not a coinage specific to the model. The phrase does not appear in <code>docs/</code> (searched, zero matches). The research index page separately claims arm’s-length internal pricing as a structural choice, but that claim lives outside this corpus.
Emergent Cooperativism	Dropped, no successor	Zero matches in <code>docs/</code> . No term occupies its role. See §5.3.
Intentional Idempotency	Idempotence, unbranded	The capitalized, branded term is gone. The underlying principle survives under a plain label as one of six infrastructure design principles in <code>docs/07-human-systems/05-synthesis-human-systems-as-infrastructure.md</code> , defined there as interactions that should be repeatable without degradation. The earlier entry’s Zen Buddhist framing did not carry forward. Note that <code>terminology.md</code> §II defines IR (Intentional Redundancy) using the same naming pattern for a different concept — overlapping roles and systems for fault tolerance, not repeatable operations.
Virtuous Cycle of Support	Dropped, no successor	Zero matches in <code>docs/</code> . The earlier entry was marked as a draft and its content paraphrases an external framework without credit. It is the weakest-sourced entry in the earlier lexicon.
Walled Garden	“walled garden”, retained as a phrase and reattached to the MEE	Survived verbatim as a phrase rather than as a defined term. <code>docs/humanized-autonomous-organization.md</code> line 4375 states that the MEE enables the walled garden without central control. The phrase originally described the ICN and Contribulo boundary specifically; it now describes the Micro Enterprise Ecosystem generally.
Mycorium	Dropped as the earlier branded term; the name reappears in chapter 09 with a different meaning	The fungal-network sense is gone. The metaphor resurfaced independently as “mycelial network” for the HAO layer and “mycorrhizal networks” for SEPs in <code>docs/humanized-autonomous-organization.md</code> and <code>docs/12-icn/01-structural-features.md</code> . Separately, the name itself is used in <code>docs/09-adversarial-analysis/</code> in a different sense; see §5.7.
Holographic Organization	Dropped as a named term; the core claim survives as “fractal”	Zero matches for “holographic” in the corpus text. The claim of self-similar structure and values recurring at every scale persists under “fractal”: fractal, polycentric architecture in <code>docs/12-icn/01-structural-features.md</code> , fractal governance layers in <code>docs/humanized-autonomous-organization.md</code> , and trust architecture described as nested and fractal in <code>docs/13-member-trust-union/01-multi-level-trust-architecture.md</code> .

Twenty-one rows. The middle column sorts each row into exactly one of three buckets, so the count below can be checked against the table directly.

- **Eleven rows name a current corpus term that stands in for the earlier one**, whether by rename, survival unchanged, absorption, or retention of the phrase: Cell, Cellular Venture, Cellular Degradation, Cellular Collapse, Cellular Lifecycle Management, Dynamic Contract, Dynamic Operating Agreement, Polycentric Governance, Polycentricity, Intentional Idempotency, and Walled Garden.
- **Three name a current term but mark the equivalence unconfirmed**: Cellular Organization, Cellular Integration, and Dynamic Cross-Network Contracts.
- **Seven name no current term at all**: Cellular Culture, Polycentric+, Arm’s Length, Emergent Cooperativism, Virtuous Cycle of Support, Mycorium, and Holographic Organization.

Cutting across those buckets, four rows record content that outlived the word carrying it. In two of them a current term stands in the earlier term’s place, so they fall in the first bucket: Idempotence is a named design principle, and “walled garden” is retained as the same phrase. In the other two the branded term has no successor term and only the underlying claim persists, in ordinary descriptive language rather than in any defined term, so they fall in the third: “fractal” for Holographic Organization, and “mycelial network” and “mycorrhizal networks” for Mycorium. §5.2, §5.4, and §5.7 treat all four.

5. Dropped and later-recovered concepts

Each entry below states what the concept meant and how much of it the current corpus covers. Some concepts were later recorded as exploratory designs in chapter 06 even though the earlier lineage draft treated them as absent. Where no source records a reason for an initial omission, this appendix does not supply one. No reason was reconstructed for any entry in this section.

5.1 Not dropped: Cellular Degradation and Cellular Collapse

These two terms were checked against the possibility that they had been dropped. They had not been. Both were renamed and formalized.

`docs/06-lifecycle/03-collapse-and-containment-protocols.md` was read directly for this appendix. It defines Enterprise Degradation as a pre-collapse state characterized by drift from core principles, reduced functionality, or social dysfunction, and specifies collapse under five trigger classes, a seven-phase lifecycle (Detection, Triage, Intervention, Confirmation, Containment, Deconstruction, Reintegration), three containment tiers, deconstruction steps covering asset inventory and equity resolution, and a post-collapse learning process. `docs/appendices/terminology.md` §VI carries both terms. The current treatment is considerably more developed than the earlier lexicon entries, which were unstructured paragraphs.

This entry is recorded to correct the premise with evidence rather than to describe a loss.

5.2 Intentional Idempotency

What it meant: processes designed so that redundant or revisited decisions do not cascade into failure, drawing on both the mathematical sense of idempotence and a stated Zen Buddhist framing of detachment from outcomes.

Reason for the drop: no recorded reason.

Current coverage: partial. “Idempotence” survives as a bare design principle, defined as interactions that should be repeatable without degradation, in `docs/07-human-systems/05-synthesis-human-systems-as-infrastructure.md`, in `docs/claude-20250429/06-human-centered-design-`

elements.md , and in docs/humanized-autonomous-organization.md line 2406. It appears there as one of six listed principles, alongside Observability, Degeneracy, Intentional Inefficiency, Versionability, and Intergenerational Capacity. The philosophical framing did not carry forward, and the concept is no longer a standalone named term. The staged extraction recorded this as one of three principles; the live corpus lists six, and the six are as named here.

5.3 Emergent Cooperativism

What it meant: that the network's cooperative structures are not fixed but evolve continuously to fit context.

Reason for the drop: no recorded reason.

Current coverage: none under another name. The general theme of adaptability recurs throughout the corpus, but it is not tied to this or any single term, and no defined term occupies the role.

5.4 Holographic Organization

What it meant: that each unit reflects the whole network's values and structure fractally, so that the microcosm mirrors the macrocosm. The source is a lexicon draft in an AI interview transcript that iterates through seven generations: an initial entry, two that expand it, three that successively simplify it, and a seventh that reverts to a formal register stated to suit an academic audience. The arc ends where it began rather than at the simplest version. It was never promoted to a standalone lexicon file.

Reason for the drop: no recorded reason for dropping the term itself.

Current coverage: the core claim survives under the word "fractal" (§4). No source documents the substitution. This is a change of word without any record of a decision to change it.

5.5 The three unit types in Cellular Lifecycle Management

What it meant: that lifecycle management, onboarding, monitoring, and support were to be tailored per unit type, across three types — Operational, Mission-Driven, and Research and Exploratory. The source specifies different treatment for each. Operational units receive heavy investment in mentorship, so that reliable processes can be passed on and dependence on specific individuals is reduced. Mission-Driven units receive ethical stress-testing alongside standard viability checks, on the stated ground that they carry the pressure to balance ethical practice against market viability. Research and Exploratory units are granted the widest latitude, offset by a mandatory cross-unit reporting obligation intended to prevent isolation and to surface harmful research directions early. Source: `~/code/site-provide-coop/provide.coop/gemini-on-clm.mdx` lines 72–74.

Historical disposition: no recorded reason for the initial omission. The source is three AI-generated passes over the framework, not a ratified unit taxonomy.

Current coverage: docs/06-lifecycle/06-lifecycle-variants-by-unit-type.md now records the three types as an exploratory layer over the uniform seven-phase UME lifecycle. It includes the source's differentiated onboarding and monitoring, Operational-unit mentorship, Mission-Driven ethical stress-testing, and Research and Exploratory cross-unit reporting. That section expressly declines to treat the older stage names as a formal one-to-one mapping onto the seven current phases or to adopt the typology as policy. The earlier statement that the types were absent was a stale cross-document search result. The label CLM was renamed to UME Lifecycle Management; the current chapter records type-specific proposals without replacing the common lifecycle.

5.6 Intentional Ephemerality

What it meant: The opening prompt fragment asked whether formally dissolving a successful Cellular Venture could become a source of learning for the network. The three subsequent Gemini response bullets elaborated a stronger model: predefined lifespans, possible wind-down at peak success, leadership assessed by resource distribution before dissolution, and a question about rotation breaks into mentorship roles. The prompt does not itself specify those mechanisms, and the AI response does not establish that they were adopted. Source: `~/code/site-provide-coop/Intentional Ephemerality.txt` (prompt fragment, line 4; Gemini response, lines 6–12).

Historical disposition: no recorded reason for the initial omission. The source remains exploratory; neither the prompt nor the AI response documents adoption as network policy.

Current coverage: docs/06-lifecycle/05-intentional-ephemerality.md now records designed finitude as an exploratory founding-stage option. It includes possible wind-down at peak success and assessment of leadership by how it distributes resources before dissolution. Rotation breaks remain an open proposal with no specified scope, timing or authority. The SEP Completion phase in docs/06-lifecycle/02-sep-lifecycle-and-governance.md is adjacent but does not itself require a predefined lifespan. The earlier statement that the sharper mechanisms were absent was a stale cross-document search result. Neither the historical source nor the current chapter establishes that any network has adopted them as policy.

5.7 Mycorium

What it meant: the whole network as a self-organizing socioeconomic system modeled on fungal networks — interconnected, resilient, and able to draw on underused resources in existing markets.

Reason for the drop: no recorded reason for dropping the branded term.

Current coverage: split, and this entry carries a correction to the staged extraction. Two things happened to the term independently.

First, the fungal metaphor itself survived without the name and without attribution back to it, reappearing as “mycelial network” for the HAO layer and as “mycorrhizal networks” for SEPs, in docs/humanized-autonomous-organization.md and docs/12-icn/01-structural-features.md . Neither use carries a caveat about the limits of biological analogy.

Second, the name “Mycorium” does appear in the current corpus, in a different sense from the earlier one. docs/09-adversarial-analysis/01-threat-taxonomy.md and docs/09-adversarial-analysis/04-external-and-technical-vectors.md use it as the sources’ term for the network’s evolving shared language and history, and identify it as the target of the threat pattern the chapter names Lexical Drift. Chapter 09 uses the term to report what its sources call something, not as a live corpus term. The staged extraction for this appendix recorded zero matches for “Mycorium” in the corpus; chapter 09 was written after that extraction. Both senses are recorded here because the sources carry both, and the corpus does not reconcile them.

One feature of the earlier entry is worth recording against both later uses. Every one of its three drafts carried a Limitations section stating that “Mycorium” was a deliberately evocative metaphor rather than a literal description of operations, and warning that overemphasizing biological analogies risks obscuring the social and economic factors that actually shape the network’s behavior. That caveat did not carry forward to either successor use.

5.8 Concepts dropped without a separate entry

Four further terms in §4 are marked dropped with no successor and are not treated at length here, because the sources record nothing about them beyond the entries themselves: Cellular Culture, Polycentric+, Arm’s Length, and Virtuous Cycle of Support. No reason is recorded for any of the four. Arm’s Length is the least consequential of them: its lexicon entry reproduced a general legal definition rather than defining anything specific to the model.

6. The founding metaphor

Sources: ~/code/site-provide-coop/provide.coop/research/ai/gemini/2024-04-01-icn-interview-1/gemini_icn-interview-1_p07r01_the-pufferfish.mdx and ...p07r02_the-pufferfish.mdx , dated 2024-04-06.

This section records a metaphor that appears in the source material. The corpus does not use it, and this appendix does not adopt it as an explanatory device. It is recorded because it is the most developed metaphor in the model’s history and because the tactics derived from it are, in the sources, treated as design guidance.

6.1 The prompt and the mapping

The metaphor originates in a prompt the author wrote. In the author’s words: “Consider this. The market is the ocean. A corporation is a shark. A non-profit is a guppy. A cooperative is a jellyfish. And Contribulo, the ICN, provide.io, and all the Cells it is composed of, is a puffer fish.”

A second prompt in the same exchange adds an anatomical mapping of three named entities plus the units. In the author’s words: “Contribulo is the skin of the puffer fish. And the mouth, and anus. provide.io is the brain, the ICN is the internal ecosystem, and Cells are the organs.”

Anatomy	Entity
Skin, mouth, and anus — the interface to the outside world	Contribulo
Brain	provide.io
Internal ecosystem	The ICN
Organs	Cells (now UMEs)

Three of the four map to entities the corpus still names. Contribulo is the boundary entity, and its role as the surface in contact with the external market corresponds to the Public Market Interface function the corpus assigns it. provide.io is a real entity. The ICN is the reference network. The fourth, Cells, is the earlier name for UMEs (§4).

6.2 Tactics derived from the metaphor

The following were generated across both response rounds of the exchange, from the pufferfish’s biology. They are summarized rather than quoted, since they are model-generated text rather than the author’s words. None was implemented as described, and none appears in the corpus.

- **Defense as a risk to its holder.** The model’s central advantage is also its central hazard if deployed carelessly, and needs safeguards and deliberate timing.
- **Inflation control.** Pufferfish inflate for a reason rather than constantly. Read as a caution against process growth for its own sake.
- **The venom factor.** Conflict between the network and an individual unit is framed as the analogue of a self-inflicted venom risk, which is the stated reason conflict-resolution mechanisms are treated as load-bearing.
- **The feint.** Some predators are deterred by the inflated form alone. Read as signaling strength while resources are low, in order to buy time rather than to deceive.
- **Near-360-degree vision.** The pufferfish’s independently moving eyes are read as an analogue for network-wide situational awareness across distributed units, conditional on investment in knowledge-sharing tooling.
- **Chain-reaction risk.** Proposes dormant units that can be activated quickly in a crisis, and stress-testing whether units can operate if cut off from central resources.
- **Fail fast, learn faster.** The structure is described as forgiving of unit-level experiments, with the caveat that lessons from failure must be distributed network-wide.

The second response turns from strategy to vocabulary management, proposing that terms such as Mycorium and Emergent Cooperativism be taught through mini-courses, an awards program, and a language-ambassadors role. That exchange is where the earlier lexicon’s most elaborate vocabulary-management proposals originate. None was implemented, and both terms it proposed teaching were later dropped (§5.3, §5.7).

6.3 The corpus is otherwise metaphor-poor

No comparable sustained, multi-part metaphor appears in docs/. The one dedicated metaphor note in the corpus, docs/notes/metaphor-family-systems.md, is short and in a different register: a few lines describing organizations that grow up as dysfunctional families under leaders missing the skills that would help them lead. It has no anatomical mapping and derives no tactics.

The corpus is not metaphor-free. Single-phrase metaphors survive independently, as recorded in §4: “walled garden”, “mycelial network”, “mycorrhizal networks”, and “fractal”. None is assembled into a governing image with an internal logic in the way the pufferfish was, and the corpus makes no argument from any of them.

HAO Open Questions Register

A record of the contradictions and unresolved questions the source material carries, checked against where the corpus itself takes a position and where it stays silent.

1. Scope, method, and the central finding

This appendix is written from two staged extractions: 04-open-questions.md, which registers 17 questions across five clusters drawn from material outside the corpus — working notes on the ICN’s entities and the MMM (Members, Mission, Market — the ordering that puts members before mission before market) working documents’ own “Still open” sections — and the ## Still open section of 05-mmm.md, which independently traces the nine MMM-proper questions to the same three MMM source files. No upstream archive was reopened to write this appendix. Every value and every source path named below was checked directly against the primary file cited, not accepted on the strength of either staging summary.

The corpus discussed throughout is the HAO (Humanized Autonomous Organization) — the network’s coordinating framework — as recorded in ~/code/hao/docs.

Central finding. The staged extraction registers 17 questions across five clusters. The active corpus now takes a substantive position on none of them. It previously asserted a 77%/23% Contribulo split for entries 2 and 3 without recorded reasoning, while the historical source material also carries a conflicting 77%/33% figure. The owner has withdrawn the active 77%/23% allocation pending a decision; this editorial correction does not reconcile the source figures or establish a replacement allocation. On the other 15 questions, the corpus either takes no position that a targeted search could locate, or reaches only adjacent material that does not answer the question as framed.

Entry 1 has since moved twice. The governance chapter initially reproduced an ownership-percentage proposal from informal notes. On 20 September 2026, the owner withdrew that proposal from the model and from generated-media source sets. Section 3.1 retains the decision record without repeating the deprecated percentage. This editorial exclusion does not establish a replacement ownership rule or a substantive model position.

One further question, present in neither staging file, is added here as §8: whether “Mycorium” — a term dropped from the corpus’s working vocabulary and independently reused by chapter 09 in a different sense — is used consistently. It was flagged for this register by the lineage appendix (docs/appendices/lineage.md §5.7), not by the open-questions staging pass, and is verified independently against its three named sources in §8. With it, this register carries 18 questions across six clusters. The count of questions on which the active corpus takes a substantive position is zero: the Mycorium question is one the corpus does not resolve either, though §8 notes a way in which it differs from the other 17.

A closing entry, §9, records a citation-integrity task first registered in docs/appendices/references.md and now completed. It is not counted among the 18 questions because it is a corpus-maintenance record rather than an owner decision or a new contradiction.

On sourcing. Every “corpus position” statement below reflects a targeted search of ~/code/hao/docs for the specific terms and figures the entry names. Where a search returns nothing, that is reported as a search result, not as proof the topic is unaddressed anywhere in the corpus under different words.

2. Summary register

#	Cluster	Question	Corpus position?
1	Ownership and control	A deprecated draft ownership-percentage proposal for incubated units	Withdrawn — see §3.1
2	Ownership and control	The Contribulo split that sums to 110% (a 77% share held in perpetuity, plus a remaining 33%)	Withdrawn — allocation unresolved
3	Ownership and control	The Contribulo perpetuity share, 77/23 vs. 77/33, and its chain of custody	Withdrawn — allocation unresolved
4	Ownership and control	Potential conflicts among incubation, service provision, and unit governance	No
5	Ownership and control	How temporary incubator controls should taper as a unit matures	No
6	Membership	Assessing a member’s contribution without it becoming a performance review	No
7	Membership	Whether there is a membership floor, and what it guarantees	No
8	Membership	How members of different units relate to each other inside a shared venture	Partially adjacent, not resolved
9	Mission	Who authors a unit’s mission	No
10	Mission	What happens when a unit’s mission is judged to have been wrong	No

#	Cluster	Question	Corpus position?
11	Mission	How much inter-unit overlap is healthy versus wasteful	No
12	Market	How long a new unit is given before its viability is assessed	No
13	Market	Who has authority to subsidize a strategic but unprofitable unit	No
14	Market	How the entity structure handles a unit that wants outside investment	No — points back to 1, 4, and 5
15	Open-source maintenance	Defining which projects count as critical	No
16	Open-source maintenance	Funding a critical project nobody commercial wants to fund	No
17	Open-source maintenance	Taking over unowned maintenance without recreating the trust problem it addresses	No
18	Vocabulary	Which sense of “Mycorium” the corpus intends	No

3. Cluster A — Ownership and control of incubated units

Five historical questions trace to informal founder notes and a set of working notes titled “Still open” in `~/code/site-provide-coop/provide.coop/docs/icn/entities.mdx`, itself flagged in its own text as reconstructed from notes that were never published. The ownership-percentage proposal formerly summarized here has been withdrawn from the active model; the original files remain private research evidence.

3.1 Entry 1 — Withdrawn ownership-percentage proposal

Historical question. Earlier informal notes assigned provide.io a fixed ownership or voting share in each incubated Cell (the earlier term for a UME — United Micro Enterprise, a small self-managing venture team of up to ~15 people). The notes disagreed about the legal form and exact threshold, and none was a ratified specification.

Decision. On 20 September 2026, the owner determined that the proposal should not appear in the HAO corpus or generated content. The active governance chapter no longer reproduces it. Canonical NotebookLM bundle construction rejects the deprecated assertion, and media that states it is ineligible for publication. The original notes remain in the private research archive so that the provenance record is not falsified by rewriting historical sources.

What remains open. The corpus has no replacement rule for ownership of an incubated unit. Any future rule requires an explicit owner decision and independent legal and cooperative-governance review; it must not be inferred from the withdrawn notes.

3.2 Entry 2 — The Contribulo split that sums to 110%

Question. Does the stated Contribulo ownership split add up to 100%?

As recorded. `entities.mdx` lines 125–127: the source states that one note gives provide.io a 77% share of Contribulo held in perpetuity, and separately refers to a remaining 33% available for participation in public markets — a total of 110% — and records that no reconciliation of the two figures was ever established.

Root source of the 77% and 33% figures, both giving the same non-summing pair: `~/code/provide.io/content/english/governance.md` lines 347 and 349, and `~/code/papr/coop-bi-003.md` lines 17 and 19, which carries the identical wording as the prompt `governance.md` was drafted from — provide.io holding 77 percent of Contribulo in perpetuity under an unequivocal Operating Agreement, with the other 33 percent available for participation on public stock exchanges.

Corpus position. A previous version of the active corpus stated 77%/23% (summing to 100) wherever it gave a Contribulo split, including the consolidated manuscript and the public-market sections of chapters 02, 04, 08, and 12. That allocation has been withdrawn from active model prose pending an owner decision. The conflicting 77%/33% source figures remain historical evidence; neither pair is an adopted allocation. The active corpus gives no reason for the earlier substitution of 23% for 33%. See entry 3 for where that substitution first enters the source chain.

3.3 Entry 3 — The Contribulo perpetuity share: 77/23 vs. 77/33, and its chain of custody

Question. The same underlying figure as entry 2, recorded separately because the source chain shows where the discrepancy entered the record, not only that it exists.

Chain of sources:

- `~/code/papr/coop-bifurcate.md` — the initiating prompt, requesting a bifurcation strategy; no percentages appear in it.
- `~/code/papr/coop-bi-003.md` — the human-authored follow-up prompt, stating the 77%-in-perpetuity, 33%-for-public-markets figures (identical wording to `governance.md`).
- `~/code/papr/bi-gpt-r-003.md` line 14 — a language model’s own restatement of that same prompt, which substitutes a different number without marking it as a change: it states that provide.io retains a 77% stake in Contribulo, with the remaining 23% available for public market participation. Nothing in the file flags the substitution or notes that the human prompt it restates gave 33.
- `~/code/provide.io/content/english/governance.md` — carries the original 77/33 figures forward into the site content as it currently stands.

Corpus position. A previous version adopted the 77/23 figure — the same number a language model substituted, unflagged, in `bi-gpt-r-003.md` years earlier — without citing that exchange or any other reason for the change. The match may be coincidental, since 23 is simply the arithmetic complement of 77, rather than a traceable editorial decision. The allocation is now withdrawn from active prose and remains unresolved pending an owner decision; both conflicting figures are retained here solely as historical source evidence.

3.4 Entry 4 — Potential conflicts among operational roles

Question. Can one entity incubate units, provide ongoing services, participate in unit governance, and direct network-level resource flows without creating a conflict of interest?

As recorded, raised and left unanswered. `entities.mdx` lines 132–136: the source states this arrangement directly and records no answer to it. It separately notes that independent arbitration appears to be the available mitigation, while stating that the question needs review by someone qualified to give a considered answer.

Corpus position. The ownership premise in the source framing was withdrawn in §3.1 and must not be treated as part of the HAO model. A narrower operational-role question remains: the corpus describes Exit Protection (`docs/11-deployment/02-provide-io-as-deployment-engine.md` line 44) and golden shares and sunset clauses (`docs/02-structure/05-public-market-interfaces.md` line 118), but it does not specify independent review, recusal, service-provider replacement, or arbitration rules for an incubator that also supplies ongoing services. No position taken on that narrower conflict-of-interest question.

3.5 Entry 5 — How temporary incubator controls should taper

Question. Which temporary controls, if any, may an incubator exercise, and how should they taper as a unit matures?

As recorded. `entities.mdx` lines 128–131 proposed heavier early control that loosens on agreed milestones and recorded no decision. Its equity premise is withdrawn under §3.1. The broader question remains relevant to non-equity controls such as approval rights, required services, technical access, or temporary governance seats.

Corpus position. The corpus specifies no equity-based incubator control. It also does not define a milestone-based handoff for the non-equity controls listed above. No position taken on that narrower question.

4. Cluster B — Membership

Three questions from `~/code/site-provide-coop/provide.coop/docs/mmm/01-members.mdx`, “Still open” section, checked directly against that file. The file carries a working-document notice: not yet reviewed.

- **Contribution assessment.** How a member’s contribution can be assessed without the process becoming a performance review. Corpus position: not addressed. The phrase “performance review” now appears once in `~/code/hao/docs`, in `docs/03-governance/03-accountability-mechanisms.md` §3.3.C, which names “Performance Review Teams” as one of five accountability bodies its MTU-sourced material states with no stated composition, appointment process, or procedure. That entry does not describe how a member’s contribution is assessed or bear on the tension this question poses between assessment and performance review; it names a body’s title and states that the source is silent beyond the title. This search no longer returns zero hits, and the result is recorded here so the register stays accurate, but the question itself remains open.
- **The membership floor.** What a member is guaranteed regardless of how much they have contributed; the source states that a floor should exist and has not been set. Corpus position: not addressed. A search for “floor” now returns one hit — `docs/13-member-trust-union/09-physical-branch-model.md` states that compensation rests on a living-wage floor for all contributors — but that floor is contribution-linked staff compensation, not a guarantee attaching to membership regardless of contribution, so it does not answer the question. The hit is recorded because it postdates this entry.
- **Cross-unit membership relations.** How members of different Cells relate to each other inside a shared Cellular Venture — a SEP (Strategic Enterprise Partnership), a joint venture between teams, in current terms. Corpus position: partially adjacent material exists — SEP governance specifies a double-consent protocol (all members plus a HAO coordination representative) for capital allocation and role changes (`docs/11-deployment/04-sep-structuring.md` line 58) — but this specifies a voting mechanism, not how membership weight or identity carries across UME boundaries inside a shared SEP. Treated as not clearly resolved, not as adopted.

5. Cluster C — Mission

Three questions from `~/code/site-provide-coop/provide.coop/docs/mmm/02-mission.mdx`, “Still open” section, checked directly against that file.

- **Mission authorship.** Whether a unit’s mission is written by the unit itself or by whoever established it.
- **Mission revision.** What happens when a unit judges its own mission to have been wrong.
- **Inter-unit overlap.** How much overlap in mission between units is healthy activity versus waste.

Corpus position, all three. `~/code/hao/docs` discusses Mission Drift Detection at length as a risk to monitor and mitigate (`docs/humanized-autonomous-organization.md` line 6352 and elsewhere), but this addresses drift away from a mission that is assumed already settled. It does not state who has authority to author or revise that mission, and it does not address the inter-unit overlap question. Not addressed.

6. Cluster D — Market viability

Three questions from `~/code/site-provide-coop/provide.coop/docs/mmm/03-market.mdx`, “Still open” section, checked directly against that file.

- **Assessment window.** How long a new unit is given before its viability is assessed.

- **Subsidy authority.** Who decides to sustain a strategically important but underperforming unit, and on what authority.
- **Outside investment.** How the entity structure handles a unit that wants outside investment. The source’s own text points this question back to the entity-ownership questions in §3 above, stating that the ownership figures there are unresolved.

Corpus position, all three. No explicit assessment window for a new UME was found in `docs/06-lifecycle/01-ume-lifecycle-and-transition-states.md`, read directly for this check. A loosely adjacent concept, time-bound probation in place of punitive bans, appears in `docs/13-member-trust-union/03-governance-and-risk-distribution.md` line 113, but it concerns sanctions against individual members inside the MTU (Member Trust Union) — the network’s credit-union-like financial institution — not a new-unit evaluation period. No decision authority for subsidizing an underperforming but strategic unit was located. Not clearly resolved on any of the three questions.

7. Cluster E — Open-source maintenance

Three questions from `~/code/site-provide-coop/provide.coop/docs/story/the-open-source-cells.mdx`, “Still open” section.

- **Defining “critical.”** Which projects count as critical, and who makes that judgment; the source states this is a judgment call with no neutral answer.
- **Funding an orphaned critical project.** What happens when a project is critical but no commercial actor wants to fund it. The source names the `xz` incident as exactly this case, and calls it the hardest to fund.
- **Taking over maintenance.** How a unit can take over maintaining a project it does not own without recreating the trust problem the arrangement exists to solve.

Corpus position, all three. `docs/08-legal/03-commons-based-ip-and-licensing.md` exists and addresses commons-based IP and licensing generally, but a targeted search of that file for “critical,” “orphan,” “unmaintained,” and “xz” returns no matches. Not addressed on these specific questions, so far as a targeted search can confirm.

8. Cluster F — Vocabulary: which sense of “Mycorium” the corpus intends

This question is present in neither staging file for this appendix. It was flagged for this register by the lineage appendix (`docs/appendices/lineage.md` §5.7) during work on that appendix, and is verified here directly against the three sources it names.

Question. “Mycorium” is used in two different senses in the material available to this corpus. Which sense, if either, the corpus intends is not stated anywhere, and the two uses are not reconciled.

Sense 1 — the earlier lexicon’s economic sense. In the dropped vocabulary recorded in `docs/appendices/lineage.md` §4 and §5.7, “Mycorium” named the whole network conceived as a self-organizing socioeconomic system modeled on fungal networks: interconnected, resilient, and able to draw on underused resources in existing markets. This sense is not used anywhere in the live corpus; a search for “Mycorium” at the time the lineage appendix’s staging extraction ran returned zero matches in `docs/`.

Sense 2 — chapter 09’s cultural and linguistic sense. `docs/09-adversarial-analysis/01-threat-taxonomy.md` line 42 and `docs/09-adversarial-analysis/04-external-and-technical-vectors.md` lines 99–101 use “Mycorium” as the sources’ term for the network’s evolving shared language and history, and identify it as the target of the threat pattern the chapter names Lexical Drift — a faction shifting the connotation of shared vocabulary over time, until the terms no longer constrain the conduct they were adopted to require. Both instances were read directly for this entry. Line 42 states that Lexical Drift’s source material names its target as the Mycorium, the sources’ term for the network’s evolving shared language and history. Lines 99–101 make the same identification and cite `docs/appendices/lineage.md` §5.7 directly as the record of the term’s two senses.

The two senses genuinely differ. Sense 1 is a structural and economic description of the network as a whole, modeled on a biological system. Sense 2 is a description of the network’s shared vocabulary and institutional memory. A network resembling a fungal network in its economic structure is a different claim from a network having a vocabulary that can drift. Neither source treats one sense as a specialization or a consequence of the other.

A nuance chapter 09 itself preserves. Chapter 09’s use of “Mycorium” reports what its own source material calls something — the text is explicit that this is the sources’ term — rather than a term chapter 09 adopts as live corpus vocabulary in its own voice. This distinguishes the chapter 09 usage from a case where the corpus had coined or endorsed the term a second time on its own authority; it reads closer to citing an external label than to redefining one.

Corpus position. No position taken. This entry differs from entries 1–17 in one respect: chapter 09 is the one place in the corpus that already names this exact tension. `04-external-and-technical-vectors.md` lines 98–99 cite `docs/appendices/lineage.md` §5.7 directly and state that the corpus does not reconcile the two senses. That self-citation flags the ambiguity; it does not resolve which sense, if either, governs when the word recurs, and no source states one.

9. Citation integrity — verification completed

When this entry was registered, `docs/appendices/references.md` §3 recorded eleven unchecked defects: the wrong author for *Slicing Pie*, two Ostrom year-and-title problems, and eight identities that the staging pass could not corroborate. None had then been checked against a primary source.

Corpus position, updated 2026-08-12. The verification task is complete. Primary publisher, journal, DOI, institutional-repository, source-control, and permanent-revision records were checked for the 48 baseline works and four supplemental citations discovered during the occurrence audit. *Slicing Pie* is now attributed to Mike Moyer; *Understanding Institutional Diversity* is dated 2005; and the *Governing the Commons* conflation is corrected to the 1990 work and merged with its duplicate audit record. Of the eight identities the staging pass could not corroborate, six remain `not verifiable` within the bounded primary-source

search and two are unsupported because official DOI or issue records positively contradict them. Their live references were removed rather than silently transferred to similar works. The full audit also found two other identities with unresolved primary records, bringing the final unresolved identity-or-access count to eight.

The wider audit also corrected additional metadata defects, removed unsupported uses, and narrowed claims that ran beyond their sources.

`docs/appendices/references.md` §3 summarizes the outcomes, while `docs/appendices/citation-audit.md` preserves the record-level evidence, occurrence verdicts, and actions. This historical entry remains outside the 18 owner questions in §2 because it records a completed corpus-maintenance task rather than an owner decision.

HAO Corpus References

The canonical bibliography for works cited by the HAO corpus, followed by the separately assembled research library and its extraction notes.

1. Scope and method

This appendix distinguishes two source sets. Section 2 is the canonical bibliography for works that remain cited in the corpus after a full citation and claim audit. Section 4 is the separately assembled 116-work research library preserved from the earlier extraction pass; §5 preserves that pass's exclusions and anomalies.

The audit froze the 48 records in the former §2 as its baseline and independently scanned the numbered chapters, `docs/humanized-autonomous-organization.md`, and `docs/appendices/lineage.md` for live author-year citations. Four additional citation candidates found in that scan received supplemental records. Local file paths, corpus cross-references, statutes, and named provenance artifacts were excluded from the citation inventory unless the corpus cited an external publication through them.

For each of the 52 audit records, the audit established or tested the bibliographic identity, rebuilt every live occurrence, identified the claim or reference-list scope attached to it, and assigned an occurrence-level verdict. Evidence came from original publications, official publisher or journal records, DOI and authoritative registry metadata, author or responsible-organization archives, and official repositories. General web search was used to locate those records, not as evidence. A similarly titled work was not substituted for a defective citation unless that different work was separately inspected and adopted; no such replacement was made in this audit.

The resulting §2 contains 39 distinct works: 36 works represented by retained or merged baseline records and three retained supplemental works. Eleven baseline records were removed, one supplemental record was removed, and the duplicate Ostrom identity B035 was merged with B032. The complete occurrence-level evidence, search boundaries, verdicts, repairs, and unresolved identity or access limits are in [Full Citation Audit](#).

The research library in §4 remains a separate 116-work extraction. Comparison by canonical title finds no overlap between its 116 titles and the 39 works in §2. Its annotations retain their original status as condensations of the staged library summaries; the full citation audit did not convert those annotations into claim-level reviews of the underlying 116 works.

2. Works cited in the corpus

Thirty-nine distinct works remain in, or were added to, the canonical corpus bibliography. Entries are ordered by the canonical first-listed author or responsible organization. The invisible audit comments connect each work to its stable record in [Full Citation Audit](#); B032 and B035 identify the same 1990 book and therefore share one entry.

- Ardener, S., & Burman, S. (Eds.). (1995). *Money-Go-Rounds: The Importance of Rotating Savings and Credit Associations for Women*. Berg.
- Bauwens, M., Kostakis, V., & Pazaitis, A. (2019). *Peer to Peer: The Commons Manifesto*. University of Westminster Press.
- Bloom, S. L., & Farragher, B. (2013). *Restoring Sanctuary: A New Operating System for Trauma-Informed Systems of Care*. Oxford University Press.
- Bollier, D., & Conaty, P. (2016). *Democratic Money and Capital for the Commons: Strategies for Transforming Neoliberal Finance Through Commons-Based Alternatives*. Commons Strategies Group in cooperation with the Heinrich Böll Foundation.
- Bollier, D., & Helfrich, S. (Eds.). (2015). *Patterns of Commoning*. Commons Strategies Group in cooperation with Off the Common Books.
- Dardot, P., & Laval, C. (2019). *Common: On Revolution in the 21st Century* (M. MacLellan, Trans.; I. Szeman, Preface). Bloomsbury Academic.
- De Filippi, P., & Wright, A. (2018). *Blockchain and the Law: The Rule of Code*. Harvard University Press.
- Fairbairn, B. (2003). *Three Strategic Concepts for the Guidance of Co-operatives: Linkage, Transparency, and Cognition*. Centre for the Study of Co-operatives, University of Saskatchewan.
- Gash, M., & Odell, K. (2013). *The Evidence-Based Story of Savings Groups: A Synthesis of Seven Randomized Control Trials*. SEEP Research, Savings-led Financial Services Working Group at the SEEP Network.
- Gibson-Graham, J. K. (2006). *The End of Capitalism (As We Knew It): A Feminist Critique of Political Economy* (new ed.). University of Minnesota Press.
- Greenfield, A. (2017). *Radical Technologies: The Design of Everyday Life*. Verso.
- Hardin, G. (1968). "The Tragedy of the Commons." *Science*, 162(3859), 1243–1248.
- Ito, J., & Howe, J. (2016). *Whiplash: How to Survive Our Faster Future*. Grand Central Publishing.

- Kahn, W. A. (1990). "Psychological Conditions of Personal Engagement and Disengagement at Work." *Academy of Management Journal*, 33(4), 692–724.
- Kelly, M. (2012). *Owning Our Future: The Emerging Ownership Revolution: Journeys to a Generative Economy*. Berrett-Koehler Publishers.
- Kostakis, V., & Bauwens, M. (2014). *Network Society and Future Scenarios for a Collaborative Economy*. Palgrave Pivot/Palgrave Macmillan.
- Lessig, L. (2006). *Code: Version 2.0* (2nd ed.). Basic Books.
- Lietaer, B., Arnsperger, C., Goerner, S., & Brunnhuber, S. (2012). *Money and Sustainability: The Missing Link*. Triarchy Press.
- McMillan, J., & Woodruff, C. (1999). "Interfirm Relationships and Informal Credit in Vietnam." *The Quarterly Journal of Economics*, 114(4), 1285–1320.
- Moyer, M. (2012). *Slicing Pie: Funding Your Business Without Funds*. Lake Shark Ventures.
- Narayanan, A., Bonneau, J., Felten, E., Miller, A., & Goldfeder, S. (2016). *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction*. Princeton University Press.
- Norman, D. A. (2013). *The Design of Everyday Things: Revised and Expanded Edition*. Basic Books.
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.
- Ostrom, E. (2005). *Understanding Institutional Diversity*. Princeton University Press.
- Ostrom, E. (2010). "Beyond Markets and States: Polycentric Governance of Complex Economic Systems." *American Economic Review*, 100(3), 641–672.
- Pentland, A., Lipton, A., & Hardjono, T. (2021). *Building the New Economy: Data as Capital*. MIT Press.
- Piketty, T. (2014). *Capital in the Twenty-First Century* (A. Goldhammer, Trans.). The Belknap Press of Harvard University Press.
- Poteete, A. R., Janssen, M. A., & Ostrom, E. (2010). *Working Together: Collective Action, the Commons, and Multiple Methods in Practice*. Princeton University Press.
- Raworth, K. (2017). *Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist*. Random House Business Books.
- Ries, E. (2011). *The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*. Crown Business.
- Schley, D. G. (2010). "Origins and Development of the Cellular Organization." In M. A. Rahim (Ed.), *Current Topics in Management: Volume 14, Organizational Behavior, Performance, and Effectiveness* (pp. 119–148). Transaction Publishers.
- Scholz, T. (2016). *Platform Cooperativism: Challenging the Corporate Sharing Economy*. Rosa Luxemburg Stiftung—New York Office.
- Scholz, T., & Schneider, N. (Eds.). (2016). *Ours to Hack and to Own: The Rise of Platform Cooperativism, a New Vision for the Future of Work and a Fairer Internet*. OR Books.
- Scott, J. C. (1998). *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press.
- Service Design Network. (2019). *Touchpoint: The Journal of Service Design*, 11(1), "Service Design for Innovation and Start-ups."
- Snow, C. C., Miles, R. E., & Allred, B. B. (2003). "Managing Leading-Edge Multinational Corporations." In B. McKern (Ed.), *Managing the Global Network Corporation* (digital ed., pp. 215–233). Routledge.
- Standing, G. (2011). *The Precariat: The New Dangerous Class*. Bloomsbury Academic.
- Weeks, K. (2011). *The Problem with Work: Feminism, Marxism, Antiwork Politics, and Postwork Imaginaries*. Duke University Press.
- Wikipedia contributors. (2016). "Cellular organizational structure." *Wikipedia*, permanent revision 756005860, 21 December 2016, 12:31:29 UTC.

3. Citation integrity audit: completed

The full audit closes the earlier staging pass's 11-item citation-integrity task. That task consisted of one author-attribution error, one year inconsistency, one title-and-year conflation, and eight identities that the staging pass could not corroborate:

- B017 was corrected from Fairfield to Mike Moyer and to the verified title *Slicing Pie: Funding Your Business Without Funds*.
- B033 was normalized to the verified 2005 publication year for *Understanding Institutional Diversity*.
- B035 was identified as a duplicate, misdated form of B032 and merged into the 1990 *Governing the Commons* entry.
- B003, B005, B006, B007, B014, B019, B031, and B036 were removed after the audit found no primary record for the cited identity within the documented search boundary, or found that the supplied metadata belonged to a different work. Related publications were not substituted silently.

The full-corpus pass also resolved defects outside that original list. B009 and B047 were removed because their cited identities could not be established. B015 identifies a real book, but its pre-audit placement did not establish a relationship to the containing section, so that occurrence and its canonical-list placement were removed. S001 was removed because Dunbar's 1992 primate study does not support an 8–15-person human work-team range. B032 and B035 now share one canonical work; no replacement work was added. The remaining metadata corrections include canonical author responsibility, edition, title, publication year, publisher, DOI, and pagination where the former forms conflicted with primary records.

The audit also assessed the claims attached to each occurrence. Supported claims were retained and citations normalized; partially supported claims were narrowed to the inspected source; broad but unattached works remain only as scoped general references; unsupported source attributions were removed; and surviving model mechanisms that exceed the sources are identified as corpus proposals. The lineage account now distinguishes the verified Wikipedia revision and the two corrected scholarly chapter identities from chronology that the available repository record cannot establish.

The 48 baseline records and four supplemental records remain in *Full Citation Audit*, including records for removed and merged citations. That appendix is the controlling record for occurrence locations, verdicts, pinpoint evidence, bounded searches, forbidden historical strings, repair actions, and unresolved identity or

access limits. Section 2 above is only the canonical bibliography of the 39 works that remain cited after those decisions.

4. Research library

116 works, assembled from the deduplicated PDF collection described in §1, grouped by the seven themes the staging extraction used, with its group counts. One annotated entry per work. The entries began as condensations of machine-generated summaries; the identities and selected descriptions for `icn-0001`, `icn-0004`, `icn-0013`, `icn-0015`, `icn-0016`, `icn-0028`, `icn-0030`, `icn-0035`, `icn-0037`, `icn-0038`, `icn-0040`, `icn-0042`, `icn-0046`, `icn-0047`, `icn-0058`, `icn-0060`, `icn-0062`, `icn-0063`, `icn-0064`, `icn-0065`, `icn-0067`, `icn-0068`, `icn-0069`, `icn-0070`, `icn-0071`, `icn-0072`, `icn-0073`, `icn-0074`, `icn-0076`, `icn-0077`, `icn-0078`, `icn-0079`, `icn-0080`, `icn-0083`, `icn-0084`, `icn-0085`, `icn-0087`, `icn-0088`, `icn-0089`, `icn-0090`, `icn-0091`, `icn-0092`, `icn-0093`, `icn-0095`, `icn-0098`, `icn-0099`, `icn-0100`, `icn-0101`, `icn-0102`, `icn-0103`, `icn-0105`, `icn-0107`, `icn-0112`, `icn-0113`, `icn-0114`, `icn-0115`, `icn-0117` and `icn-0118` were later corrected against original documents. Those bounded corrections do not verify every statement in those entries or the remaining staged annotations. Entries marked `[metadata incomplete]` still lack a reliable author or publication year (5 of 116). Each entry closes with its library reference key (for example, `[icn-0013]`) for traceability back to the canonical collection. Private review notes preserve the corrections and their limits.

The subsequent checks of `icn-0012`, `icn-0014`, `icn-0017`, `icn-0021`, `icn-0022`, `icn-0026`, `icn-0027`, `icn-0029`, `icn-0034`, `icn-0036`, `icn-0039`, `icn-0043`, `icn-0061`, `icn-0066`, `icn-0094`, `icn-0097`, `icn-0106`, `icn-0108`, `icn-0109`, `icn-0110` and `icn-0111` likewise correct their source types and the limits of their staged claims.

Polycentric Governance & Commons (22)

- **Costs and Challenges of Polycentric Governance** — Michael McGinnis (revised 2011). This working paper proposes a framework for comparing governance costs. It describes participant time and effort, pressures towards standardisation, and the risk of treating polycentric order as automatically socially optimal. These are analytical arguments, not tests of HAO performance. Bounded claim review (review retained privately). `[icn-0013]`
- **Multi-Stakeholder Governance in Cooperative Organizations: Toward a New Framework for Research?** — Catherine Leviten-Reid; Brett Fairbairn (2011). This review challenges theoretical predictions of failure from high decision-making costs using limited survey and case-study evidence. It proposes common-pool governance as a research comparison, without testing that analogy or HAO's governance design. Bounded claim review (review retained privately). `[icn-0015]`
- **Redefining Polycentricity** — Chen Zhong (2014). This short author blog post concerns urban spatial clusters and their measurement, not polycentric governance among decision-making centres. One generated claim came from an unrelated "Recent Posts" sidebar. Bounded claim review (review retained privately). `[icn-0058]`
- **What is Dynamic Governance?** — Becky Bowen, J.D. (PDF created August 2014; publication date not printed). This first-person explainer describes consent for policy decisions, semi-autonomous circles, double-linking and consent elections, then suggests meeting practices. Its only reported result is the author's impression of one facilitation; it does not measure organisational effects. Bounded primary-PDF spot check (review retained privately). `[metadata incomplete]` `[icn-0026]`
- **Dynamic Governance in Theory and Application, Part I** — David L. Markell; Robert L. Glicksman (2016). *Arizona Law Review*, 58(3), 563–632. The library PDF is a pre-publication draft. This conceptual legal analysis proposes an actors–mechanisms–tools framework and examines EPA's Next Generation Compliance initiative. Its use of "dynamic governance" concerns regulatory design, not the consent-based circle method described in `icn-0026`; it does not evaluate HAO. Bounded primary-PDF spot check (review retained privately). `[icn-0029]`
- **Polycentric Governance in Theory and Practice: Dimensions of Aspiration and Practical Limitations** — Michael D. McGinnis (2016 draft; posted to SSRN in 2021). The author proposes a structure–process–outcome definition and six possible failure modes, while distinguishing the normative ideal from practical approximations. Bounded claim review (review retained privately). `[icn-0030]`
- **Distributed Governance** — Carla L. Reyes; Nizan Geslevich Packin; Benjamin P. Edwards (2017). *William & Mary Law Review Online*, 59, Article 1. This legal analysis examines governance opportunities and risks for entities using distributed ledger technology; it is not a general account of polycentric human organisations. The audit handoff misidentified the authors and work; see the primary-PDF spot check (review retained privately). `[icn-0080]`
- **Polycentric Systems of Governance: A Theoretical Model for the Commons** — Keith Carlisle; Rebecca L. Gruby (2019 journal volume; first online 2017). This literature synthesis proposes defining attributes and context-dependent enabling conditions for commons governance. It does not establish that polycentric systems consistently outperform alternatives. Bounded claim review (review retained privately). `[icn-0035]`
- **Polycentric Governance Processes** — *Author unknown* (2018). The material contains a section titled 'Part III: Polycentric Governance Processes'. `[metadata incomplete]` `[icn-0081]`
- **United People: Designing a New Model of Global Governance** — A. Georges L. Romme; Christopher Ansell; John Buck; Younghoon Choi; Rob van der Eyden; Verónica Figueroa Huencho; Edwin John; Tracy Kunkler; Johanna Mair; Albert Meijer; Renate Meyer; Karen Stephenson; Liisa Välikangas; Nathaniel Whitestone; Cordelia Yu (2018). *Journal of Asian Scientific Research*, 8(4), 152–170. This idealised global-governance design proposes small linked councils, informed-consent policy decisions and an executive board; it does not test their effects. The staged summary invents an algorithm and allocation metrics absent from the article. Bounded primary-PDF spot check (review retained privately). `[icn-0036]`
- **An Introduction to Polycentricity and Governance** — Mark Stephan; Graham Marshall; Michael McGinnis (2019). This conceptual chapter distinguishes polycentricity, polycentric arrangements and governance systems. It recounts biological uses of the term but does not establish that those uses transmitted into the 1961 Ostrom–Tiebout–Warren account. Its performance comparison is limited to some settings. Bounded claim review (review retained privately). `[icn-0042]`

- **It's Ordered Chaos: What Really Makes Polycentrism Work** — Maria Koinova; Maryam Zarnegar Deloffre; Frank Gadinger; Zeynep Sahin Mencutek; Jan Aart Scholte; Jens Steffek (2021). *International Studies Review*, 23(4), 1988–2018. This six-author theoretical forum proposes several relational lenses for interpreting polycentric global governance, illustrated in different policy fields. It does not compare their performance or establish HAO ancestry. Bounded claim review (review retained privately). [icn-0044]
- **Polycentric Collaborative Governance, Sustainable Development and the Ecological Resilience of Elevator Safety: Evidence from a Structural Equation Model** — Rijia Ding; Chongbao Ren; Suli Hao; Qi Lan; Mingbo Tan (2022). *Sustainability*, 14(12), 7124. The authors model associations between a proposed multi-actor governance framework and elevator-safety resilience using questionnaire responses from Chinese elevator-safety units. The study's sector and survey design limit transfer to HAO governance. Bounded primary-PDF spot check (review retained privately). [icn-0047]
- **The Evolution of Polycentric Governance in the Galapagos Small-Scale Fishing Sector** — Renato Cáceres; Jeremy Pittman; Mauricio Castrejón; Peter Deadman (2022). This study uses network, interview and document evidence to explore collaboration in the fishing sector. Its 28 candidate nodes and a separate 43-organisation prior network must not be conflated. Bounded handoff spot check (review retained privately). [icn-0046]
- **A Progressive Web3: From Social Coproduction to Digital Polycentric Governance** — Quinn DuPont (2023 draft, revised 8 February). This conceptual essay proposes studying decentralised autonomous organisations (DAOs) as possible sites of digital polycentric governance. Its DAO-data research programme is prospective; it does not test governance outcomes or establish HAO ancestry. Bounded claim review (review retained privately). [icn-0050]
- **Building Blocks of Polycentric Governance** — Tiffany H. Morrison; Örjan Bodin; Graeme S. Cumming; Mark Lubell; Ralf Seppelt; Tim Seppelt; Christopher M. Weible (2023). *Policy Studies Journal*, 51, 475–499. The authors test a three-mode network-motif method on a reconstructed Great Barrier Reef governance case at three time points. This is a proof of concept, not evidence that the method or polycentric governance improves HAO outcomes. Bounded claim review (review retained privately). [icn-0049]
- **Humanistic Governance in Democratic Organizations: The Cooperative Difference** — Sonja Novković; Karen Miner; Cian McMahon (eds., 2023). This edited volume examines cooperative governance across different member relationships and organisational contexts. The series' founding date is not a description of the book's findings. Bounded primary-PDF spot check (review retained privately). [icn-0089]
- **Polycentric Governing and Polycentric Governance** — Andreas Thiel (2023). This conceptual chapter distinguishes descriptions of multi-actor governance from normative criteria such as voice, exit and self-organisation. Its examples and cited frameworks do not establish HAO ancestry or performance. Bounded claim review (review retained privately). [icn-0087]
- **Conclusions: Critical Reflections on the Epistemic Adequacy of the Western Legal Approach to Square the Circle and Grant a Common Future for All** — Margherita Paola Poto (2023). In *Sustainability through Participation?*, 491–519. This concluding chapter synthesises the volume and argues for re-examining anthropocentric assumptions in environmental law through plural, participatory approaches. Its conceptual critique is not an evaluation of HAO governance. Bounded claim review (review retained privately). [icn-0051]
- **The Exchange Theory of Web3 Governance** — Darcy W. E. Allen; Chris Berg; Aaron M. Lane; Trent MacDonald; Jason Potts (2023). *Kyklos*, 76, 659–675. This conceptual article contrasts a community-voting account of Web3 governance with an exchange-based account, illustrated by Curve, Lido and Metagov. Its proposed benefits of governance-rights exchange and polycentricity are arguments, not measured outcomes or evidence of HAO ancestry. Bounded claim review (review retained privately). [icn-0053]
- **The Black Box of Power in Polycentric Environmental Governance** — T.H. Morrison; W.N. Adger; K. Brown; M.C. Lemos; D. Huitema; J. Phelps; L. Evans; P. Cohen; A.M. Song; R. Turner; T. Quinn; T.P. Hughes (2019). *Global Environmental Change*, 57, 101934. This Perspective examines how power shapes rule-setting, issue construction and implementation in polycentric environmental governance. Three staged claims were drawn from its reference list rather than its analysis; see the primary-PDF spot check (review retained privately). [icn-0056]
- **Dynamic Governance Summary** — Glenda Mattison (created); John Buck, Tena M. O'Rear and Monika Megyesi (edited; sheet dated March 2009, © 2008 Sociocratisch Centrum). This two-page reference distinguishes proposal formation, consent decisions, elections, circle meetings and leader-led operational meetings. Three staged claims join text from different columns; see the bounded primary-PDF spot check (review retained privately). [icn-0001]

Cybernetics, VSM & Holonic Systems (incl. Emergent Cooperation / Multi-Agent Models) (26)

- **Hierarchical Control System Emulation: User's Manual** — Cita M. Furlani, ed. (1985). National Bureau of Standards, NBS IR 85-3156. This manual documents a real-time emulation and simulation tool for modular hierarchical control systems, especially automated manufacturing; its introduction treats simulation as a design aid for feedback control. Bounded primary-PDF spot check (review retained privately). [icn-0004]
- **Holonic Manufacturing Systems: Initial Architecture and Standards Directions** — James H. Christensen (presented 1 December 1994, First European Conference on Holonic Manufacturing Systems, Hannover). This conference paper proposes an initial manufacturing architecture of autonomous, cooperative holons, interface functions and standards tasks. The author calls it his synthesis rather than a consortium consensus; its projected manufacturing benefits are not HAO outcomes. Bounded claim review (review retained privately). [icn-0068]
- **A Complex Adaptive Systems Model of Organization Change** — Kevin J. Dooley (1997). *Nonlinear Dynamics, Psychology, and Life Sciences*, 1(1), 69–97. This literature-based paper proposes a model of organisational change drawing on systems theory, population ecology and information processing. Its change strategies and fitness curve are conceptual, not measured HAO outcomes. Bounded claim review (review retained privately). [icn-0006]
- **Conflict Detection and Resolution in the Presence of Prediction Error** — Heinz Erzberger; Russell A. Paielli; Douglas R. Isaacson; Michelle M. Eshow (1997). Prepared for the 1st USA/Europe Air Traffic Management R&D Seminar, this paper describes the aviation-specific Conflict Probe and reports search-algorithm performance tests. It does not evaluate HAO conflict governance or the FAST feedback acronym used elsewhere in the library. Primary-PDF spot check (review retained privately). [icn-0097]
- **Holonic Manufacturing Scheduling: Architecture, Cooperation Mechanism, and Implementation** — Ling Gou; Peter B. Luh; Yuji Kyoya (1998). *Computers in Industry*, 37(3), 213–231. This manufacturing-scheduling study tests a two-level holonic optimisation method using factory-derived data on a single workstation. It does not test a distributed physical implementation or HAO governance. Five staged rows contain column-spliced text or reference/biography material; see the bounded claim review (review retained privately). [icn-0069]

- **Knowledge Workers as Fractals in a Complex Adaptive Organization** — Snunith Shoham; Alon Hasgall (2005). *Knowledge and Process Management* 12(3), 225–236. A qualitative study of interviews and observations at six organisations examines employees as knowledge-bearing subsystems and the coordination of their work; its findings do not establish HAO’s ancestry or performance. The local copy is an uncorrected proof, whose provisional author line differs from the publisher-deposited citation. Bounded claim review (review retained privately). [icn-0009]
- **Extracting the Hierarchical Organization of Complex Systems** — Marta Sales-Pardo; Roger Guimerà; André A. Moreira; Luís A. Nunes Amaral (2007). *PNAS* 104(39), 15224–15229. The paper tests an unsupervised method for identifying nested modules in simulated and selected empirical networks; it does not test organisational governance or HAO architecture. A published correction replaces Fig. 1B without changing the authors’ conclusions. Three staged quotations splice the PDF’s columns; see the bounded claim review (review retained privately). [icn-0010]
- **Turn Your Company Outside-In! How to Create the Adaptive Network Organization: A Concept Paper on Cell Structure Design, Part I** — Gebhard Borck; Valeria Junqueira; Niels Pflaeging; Andreas Zeuch (October 2008; revised February 2009). This BBTN Associates white paper proposes cell-based network design; it does not measure transformation outcomes. Three generated claims come from service-promotional backmatter and are not research evidence. See the primary-PDF claim review (review retained privately) and earlier identity spot check (review retained privately). [icn-0099]
- **The Viable System Model and Its Application to Complex Organizations** — Allenna Leonard (2009). *Systemic Practice and Action Research*, 22, 223–233. Leonard explains Stafford Beer’s five-system model, its communication channels and recursion, with a small-business example and an account of the Chilean Cybersyn project. This is Leonard’s interpretation and historical account, not an original Beer publication or evidence that the model shaped HAO. Bounded primary-PDF and claim review (review retained privately). [icn-0071]
- **The Origins and Purposes of Several Traditions in Systems Theory and Cybernetics** — Stuart A. Umpleby; Eric B. Dent (1999; online posting 2010). *Cybernetics & Systems*, 30(2), 79–103. This historical review distinguishes several systems and cybernetics traditions; it is secondary historical analysis, not evidence of HAO ancestry. One staged quotation has an OCR artifact; see the primary-PDF spot check (review retained privately). [icn-0073]
- **New Federated Collaborative Networked Organization Model (FCNOM)** — Morcoux M. Yassa; Hesham A. Hassan; Fatma A. Omara (2012). *International Journal of Cloud Computing and Services Science*, 1(1), 1–10. This paper proposes a model of collaborative networked organisations and illustrates it with an insurance case; the local PDF contains only printed pp. 3–10. Its model description does not establish measured effectiveness. Primary-PDF spot check (review retained privately). [icn-0100]
- **Fractal Organization Theory** — Janna Raye (2012). This conceptual essay advocates shared purpose, distributed decisions and two-way communication using fractal and living-systems metaphors. It does not establish organisational performance effects or HAO ancestry. Primary-PDF spot check (review retained privately). [icn-0101]
- **From Emergent Cooperation to Contextual Trust, and to General Trust: Overlapping Meso-Sized Interaction Arenas and Cooperation Platforms as a Foundation of Pro-Social Behavior** — Wolfram Elsner; Henning Schwardt (online 2014; 2015 journal issue). *Forum for Social Economics*, 44(1), 69–86. This conceptual analysis proposes that cooperative patterns in smaller arenas may support wider trust, while stressing that transfer across arenas is not determined. It does not test HAO outcomes. Primary-PDF spot check (review retained privately). [icn-0108]
- **Hierarchical Emulation & Data Assimilation into the Sediment Transport Model** — Nugzar Margvelashvili; Eddy Campbell; Laurence Murray; Emlyn Jones (2014). *Procedia Computer Science* 29, 2121–2126. In one idealised-estuary simulation, a hierarchy reflecting dependencies among sediment-model variables improves the emulator’s fit and inferred suspended-solids estimates relative to the compared non-hierarchical emulator. The study uses synthetic observations and does not test organisational hierarchy or HAO. Bounded primary-PDF claim review (review retained privately). [icn-0023]
- **Holonic Manufacturing System** — L. S. Nikam; H. K. Raval; J. S. Bagi (May 2015 per an author-uploaded record; the PDF is undated). This overview presents holonic manufacturing as a proposed response to changing production demands. Its definitions of autonomy, cooperation and integration across manufacturing activities are attributed to the 1994 HMS consortium, not independently established here; its claimed benefits are not measured outcomes. Bounded primary-PDF and claim review (review retained privately). [icn-0077]
- **Hierarchical Emulation: A Method for Modeling and Comparing Nested Simulators** — Rachel H. Oughton; Peter S. Craig (2016). *SIAM/ASA Journal on Uncertainty Quantification* 4(1), 495–519. The authors develop a Bayesian method for related versions of a computer simulator and compare its predictions with standard emulators using an ocean carbon-cycle model. The reported advantage is conditional on that experimental setup and does not test organisational hierarchy or HAO. Bounded primary-PDF claim review (review retained privately). [icn-0031]
- **Mobility Can Promote the Evolution of Cooperation via Emergent Self-Assortment Dynamics** — Jaideep Joshi; Iain D. Couzin; Simon A. Levin; Vishwesh Guttal (2017). *PLOS Computational Biology*, 13(9), e1005732. Spatially explicit evolutionary simulations and an analytical model show conditional cooperation-cohesion dynamics in mobile biological populations; they do not test human organisations or HAO. Primary-PDF spot check (review retained privately). [icn-0109]
- **Emergent Coordination Through Competition** — Siqi Liu; Guy Lever; Josh Merel; Saran Tunyasuvunakool; Nicolas Heess; Thore Graepel (ICLR 2019). In a simulated multi-agent soccer environment, the authors report coordination among reinforcement-learning agents trained with population-based methods. The findings concern that training setup, not human organisations. Primary-PDF spot check (review retained privately). [icn-0110]
- **Evolution of Holonic Control Architectures Towards Industry 4.0: A Short Overview** — Olivier Cardin; William Derigent; Damien Trentesaux (2018). *IFAC-PapersOnLine*, 51(11), 1243–1248. This literature overview treats overconnectivity and data management as possible enablers of flexible manufacturing control. Its accounts of PROSA and POLLUX relay earlier architectures; the paper does not itself test their performance or HAO governance. Bounded primary-PDF and claim review (review retained privately). [icn-0083]
- **Effective Assimilation of Technological Innovation in an Organization Characterized as a Complex Adaptive System** — Alon E. Hasgall; Niv Ahituv; Nily Naveh (2019). *Journal of Innovation Management* 7(2), 38–58. A questionnaire study of 300 employees reports associations between complex-adaptive-system characteristics and attitudes towards technological change. The published PDF conflicts on whether the sample covered 15 or 11 organisations, and its regression description conflicts with its stated variable direction. It does not test HAO or establish a causal effect. Bounded primary-PDF review (review retained privately). [icn-0041]
- **Emergent Cooperation Through Mutual Information Maximization** — Santiago Cuervo; Marco Alzate (2020). *arXiv:2006.11769v1*. This preprint proposes a decentralised reinforcement-learning method and compares it with a baseline in one simulated commons game. Its result does not establish cooperation

across all social dilemmas or in HAO. Primary-PDF spot check (review retained privately). [icn-0111]

- **A Sustainable Ecosystem through Emergent Cooperation in Multi-Agent Reinforcement Learning** — Fabian Ritz; Daniel Ratke; Thomy Phan; Lenz Belzner; Claudia Linnhoff-Popien (2021). *ALIFE 2021*, article 74. The paper reports results from a simulated predator-prey environment; the handoff's ICRA attribution is incorrect. See the primary-PDF spot check (review retained privately). [icn-0112]
- **Emergent Cooperation from Mutual Acknowledgment Exchange** — Thomy Phan; Felix Sommer; Philipp Altmann; Fabian Ritz; Lenz Belzner; Claudia Linnhoff-Popien (2022). *AAMAS 2022*, pp. 1047–1055. The paper evaluates Mutual Acknowledgment Token Exchange (MATE) in three simulated social-dilemma domains. Two staged claims are bibliography entries, and another joins unrelated columns; see the primary-PDF spot check (review retained privately). [icn-0113]
- **Emergent Cooperation and Deception in Public Good Games** — Nicole Orzan; Erman Acar; Davide Grossi; Roxana Rădulescu (2023). *Adaptive and Learning Agents Workshop (ALA 2023)*. The paper studies communication among reinforcement-learning agents in an extended public-goods game. The audit handoff omitted an author and gave the wrong venue; see the primary-PDF spot check (review retained privately). [icn-0114]
- **Emergent Cooperation under Uncertain Incentive Alignment** — Nicole Orzan; Erman Acar; Davide Grossi; Roxana Rădulescu (2024). *AAMAS 2024*. The paper studies reputation and intrinsic reward in simulated public-goods games with uncertain incentive alignment. The staged title and several generated quotations are incorrect; see the primary-PDF spot check (review retained privately). [icn-0115]
- **Holographic Organizations: An Introduction** — Institute for the Advancement of Service, Susan's Blog (publication date and individual byline unconfirmed; PDF captured in 2024). This institutional teaching proposes that members and activities express a shared organisational philosophy while retaining distinct perspectives; it also recommends using both holographic and hierarchical elements as appropriate. Bounded primary-PDF spot check (review retained privately). [metadata incomplete] [icn-0067]

Cooperative Economics & Worker Ownership (10)

- **“Pay Equity” through “Equitable Payment”** — Alan Pearson (1993 revision of a 1992 Festschrift article). Pearson proposes Time-Span of Discretion as a possible measure for equal pay for work of equal value and reports a 1986 pilot involving 13 employees in one Montréal unit. The small, gender-uneven sample and the paper's call for further research do not establish a general pay-setting rule or validate HAO compensation. The canonical filename incorrectly says “equitable profit”; see the primary-PDF claim review (review retained privately). [icn-0005]
- **Can Redistribution Accelerate Growth and Development?** — François Bourguignon (June 2000 first draft). Prepared for ABCDE/Europe, this literature synthesis examines conditions under which redistribution of productive assets might address market imperfections and support growth. It reports no new empirical results, finds aggregate cross-country evidence inconclusive, and does not show that redistribution in general accelerates growth or validate HAO's trickle-out mechanism. The research-library cover says January 2000, but the manuscript itself says June; see the primary-PDF claim review (review retained privately). [icn-0007]
- **Solidarity as a Business Model: A Multi-Stakeholder Cooperatives Manual** — Margaret Lund, primary author (2011). Published by the Cooperative Development Center at Kent State University, a programme of the Ohio Employee Ownership Center. This practice manual discusses member classes, allocation of governance rights, distribution of surplus, dissolution and transfer rights, and case examples; it presents options rather than a tested rule for HAO. Bounded primary-PDF spot check (review retained privately). [icn-0014]
- **Providing the Right Support: Are the Training Needs of Multi-Stakeholder Cooperatives Unique?** — Courtney Berner (October 2013). University of Wisconsin Center for Cooperatives informal interview study of seven US multi-stakeholder cooperatives, six in food or agriculture. It concludes that most board-training needs resemble those of other cooperatives, while recommending attention to governance structure, shared mission, and communication and conflict skills. The small, sector-heavy sample does not establish general effects of those recommendations. Bounded primary-PDF spot check (review retained privately). [icn-0022]
- **Building Shared Entrepreneurship** — Joe Rinehart (2015). Democracy at Work Institute guide for small-business support organisations. It introduces worker-cooperative development questions and gives short case examples; its general claims about job creation and business survival are not independently tested in this brochure. The publisher's resource page supplies the byline absent from the PDF. Bounded primary-source spot check (review retained privately). [icn-0027]
- **Limited Liability Companies as Worker Cooperatives** — Sarah Sexton (2009, hosting resource record). An Insight Center guide presenting LLC operating-agreement choices for worker cooperatives, including management, voting and withdrawal. Its absolute limited-liability wording has exceptions in the same document; its California filing and tax details are historical, not current legal advice. See the primary-PDF claim review (review retained privately). [icn-0037]
- **The Spectrum of Community Engagement to Ownership** — Rosa González of Facilitating Power (2019; developed in partnership with Movement Strategy Center). This practitioner tool places six stances from ignoring a community through deferring to community ownership alongside assessment and planning exercises. Its resource-allocation percentages are presented without an empirical derivation, and its claimed benefits are not outcome findings. Bounded primary-PDF spot check (review retained privately). [icn-0039]
- **Building “Next Generation” Democratic Workplaces to Reduce Inequality and Empower Workers: Evidence and Policy Implications from Buffalo-Niagara** — Russell Weaver (October 2020). Cornell ILR Buffalo Co-Lab policy report. Its mixed-methods Buffalo-Niagara case study finds associations between union or ESOP presence and wage measures, using census-tract comparisons and modelled employee distributions; the author explicitly says the data do not support causal conclusions. It also reviews worker cooperatives and social-mission businesses and proposes policy measures. Bounded primary-PDF spot check (review retained privately). [icn-0043]
- **The Slow Demise of Loconomics** — Danny Spitzberg (2021). Interview with Loconomics co-founder Joshua Danielson, first published in *STIR* and reposted by Grassroots Economic Organizing. Danielson recounts the platform's product pivot, funding constraints and closure; his account is not an independently verified performance study. Bounded source spot check (review retained privately). [icn-0092]

- **Cooperative Housing Toolbox: A Practical Guide for Cooperative Success** — Northcountry Cooperative Foundation, in partnership with Northcountry Cooperative Development Fund (© 2003). This housing-cooperative guide covers member rights and responsibilities, housing equity forms and board governance. Its PDF creation date is in 2004; that metadata is not a publication date. Bounded primary-PDF spot check (review retained privately). [icn-0118]

Organizational Psychology & Human Systems (25)

- **Self-Determination Theory and the Facilitation of Intrinsic Motivation, Social Development, and Well-Being** — Richard M. Ryan; Edward L. Deci (2000). *American Psychologist*, 55(1), 68–78. The article examines conditions affecting intrinsic motivation and psychological-need satisfaction. The generated OIT name and audit handoff's claim-ID mapping are unreliable; see the primary-PDF spot check (review retained privately). [icn-0098]
- **Organizational Behavior: Human Behavior at Work** — John W. Newstrom (12th ed., 2007). McGraw-Hill/Irwin. This textbook introduces organisational behaviour across individual, group and whole-system levels. Its definitions and stated goals are pedagogical framing, not evidence that a particular HAO practice works. Bounded primary-PDF spot check (review retained privately). [icn-0070]
- **Perceived Match or Mismatch on the Gottman Conflict Styles: Associations with Relationship Outcome Variables** — Dean M. Busby; Thomas B. Holman (2009). *Family Process*, 48(4), 531–545. In a self-report sample of 1,983 heterosexual couples, perceived conflict-style categories were associated with relationship satisfaction, stability, problems and stonewalling. The sample is nonrepresentative and the analysis cross-sectional; it establishes neither causal effects nor transfer to HAO teams. Bounded primary-PDF spot check (review retained privately). [icn-0012]
- **Universal Mental Health Program: An Extension of Life Skills Education to Promote Child Mental Health** — Narayana Manjunatha; Sahoo Saddichha (2011). *Indian Journal of Psychiatry*, 53(1), 77–78. This letter to the editor proposes that WHO promote life-skills education within a universal mental-health programme; it does not report an implemented programme or measure its effects. Bounded primary-PDF spot check (review retained privately). [icn-0076]
- **Corporate Reputation, Identity and Trust** — Wesselina Andria Johanna van der Merwe (chapter 5 of a 2013 University of Pretoria PhD thesis, *Towards a Conceptual Model of the Relationship Between Corporate Trust and Corporate Reputation*). This theoretical chapter proposes that trustworthy conduct and an ethical corporate identity inform stakeholder assessments of reputation, with trust treated as an outcome. The earlier attribution to Fombrun and Van Riel, whom the chapter cites, was incorrect; the chapter does not test the proposed model or HAO member trust. See the primary-PDF claim review (review retained privately). [icn-0003]
- **Does Reputation Work to Discipline Corporate Misconduct?** — Jonathan M. Karpoff (2012), chapter 18 in *The Oxford Handbook of Corporate Reputation*, pp. 361–382. This review of empirical studies reports substantial estimated reputation losses for some misconduct affecting business counterparties, but negligible average losses in the environmental-violation studies discussed. The earlier research-library title was wrong; the chapter does not establish a general self-enforcing substitute for HAO governance. See the primary-PDF claim review (review retained privately). [icn-0018]
- **The Building Blocks of Corporate Reputation: Definitions, Antecedents, Consequences** — Charles J. Fombrun (2012), chapter 5 in *The Oxford Handbook of Corporate Reputation*, pp. 94–113. This theoretical synthesis compares seven conceptual lenses and proposes an integrative reputation framework for future empirical testing; it does not measure the framework's effects or test an HAO intervention. The generated claim set reuses one raw claim ID for two distinct statements. See the primary-PDF claim review (review retained privately). [icn-0019]
- **Predicting Future Conflict Between Team-Members with Parameter-Free Models of Social Networks** — Núria Rovira-Asenjo; Tània Gumí; Marta Sales-Pardo; Roger Guimerà (2013). *Scientific Reports*, 3, article 1999. In 16 chemical-engineering student teams, a network model ranked later changes in willingness to work together above a within-team null model; a structural-balance comparator did not. The survey proxy is narrower than observed conflict, and the study does not validate deployment in HAO teams. Bounded primary-PDF spot check (review retained privately). [icn-0021]
- **Psychosocial Deprivation, Executive Functions, and the Emergence of Socio-Emotional Behavior Problems** — Jennifer Martin McDermott; Sonya Troller-Renfree; Ross Vanderwert; Charles A. Nelson; Charles H. Zeanah; Nathan A. Fox (2013). *Frontiers in Human Neuroscience*, 7, article 167. This follow-up analysis of the Bucharest Early Intervention Project examines flanker-task performance and error-monitoring measures in eight-year-old children. The abstract's foster-care versus care-as-usual Pe claim is not matched by a reported direct significant comparison in the results; other reported associations are outcome- and group-specific. The study does not test HAO teams or adult member assessment. See the primary-PDF claim review (review retained privately). [icn-0020]
- **Hebbian Learning and Predictive Mirror Neurons for Actions, Sensations and Emotions** — Christian Keysers; Valeria Gazzola (2014). *Philosophical Transactions of the Royal Society B*, 369, article 20130175. Labelled an opinion piece, it proposes a Hebbian account of how predictive mirror-like responses might develop and draws on earlier experimental work; it does not directly test that developmental account or HAO coordination. Two generated evidence snippets splice separate PDF columns, and the claim set remains quarantined. See the primary-PDF claim review (review retained privately). [icn-0024]
- **Stable Conflict Resolution Styles and Commitment: Their Roles in Marital Relationship Self-Regulation** — Rebecca Suzanne Boyd (2014). Brigham Young University master's thesis. Using cross-sectional self-reports from 2,262 first-married individuals, it examines associations among self-classified conflict style, commitment and relationship self-regulation. The earlier bibliography's authors were wrong. The study does not test workplace teams or HAO conflict procedures; its claimed validating-style advantage is narrower than the abstract suggests. See the primary-PDF claim review (review retained privately). [icn-0025]
- **Cultures of Resilience: Ideas** — Ezio Manzini and Jeremy Till, eds. (2015). London: Hato Press. ISBN 978-1-910239-10-0. This University of the Arts London project volume contains a shared base text on socio-technical resilience and separately authored contributions; the base text discusses diversity, redundancy, feedback and experimentation as possible characteristics of resilient systems. Bounded primary-PDF spot check (review retained privately). [icn-0028]
- **To Have Control Over or to Be Free From Others? The Desire for Power Reflects a Need for Autonomy** — Joris Lammers; Janka I. Stoker; Floor Rink; Adam D. Galinsky (2016). *Personality and Social Psychology Bulletin*, 42(4), 498–512. Across nine studies, the authors report that autonomy, more than influence over others, is associated with the desire for power. This individual-level finding does not test HAO governance. Bounded primary-PDF spot check (review retained privately). [icn-0078]

- **Hierarchical Rank and Principled Dissent: How Holding Higher Rank Suppresses Objection to Unethical Practices** — Jessica A. Kennedy; Cameron Anderson (2017). *Organizational Behavior and Human Decision Processes*, 139, 30–49. The authors report lower principled dissent with higher rank across an archival survey and two experiments; the canonical PDF is a publisher preview without the full study methods or results. Bounded primary-PDF spot check (review retained privately). [icn-0034]
- **The Psychology of Harmony and Harmonization: Advancing the Perspectives for the Psychology of Sustainability and Sustainable Development** — Annamaria Di Fabio; Akira Tsuda (2018). *Sustainability*, 10(12), article 4726. This perspective proposes a psychological account of harmonisation involving oneself, others and the natural world; it is not an empirical evaluation of HAO groups or institutions. Bounded primary-PDF spot check (review retained privately). [icn-0038]
- **Organizational Behavior** — OpenStax, Rice University (2019). The archived PDF is this OpenStax textbook, not the differently titled work suggested by its canonical filename. Its copyright page declares CC BY 4.0 for that PDF; the current OpenStax web edition displays different reuse terms, so rights must be assessed by version before reuse. Bounded primary-PDF spot check (review retained privately). [icn-0084]
- **The ‘Virtuous Cycle’: How to Create Compounding Successes** — Shawn Achor (published 23 September 2019; updated 1 October 2019). *SUCCESS* magazine essay. Achor proposes reciprocal support and recognition as a “virtuous cycle” and illustrates the idea with hypothetical workplace examples. The archived article says research supports a happiness-before-success direction but identifies no study, methods or data; it is not empirical evidence for that claim or for HAO outcomes. Bounded primary-PDF claim review (review retained privately). [icn-0059]
- **Turning Vicious Cycles Into Virtuous Ones: the Potential for Schools to Improve the Life Course** — Mitchell D. Wong; Karen Hunter Quartz; Marisa Saunders; Ben P. L. Meza; Saltanat Childress; Teresa E. Seeman; Rebecca N. Dudovitz (May 2022). *Pediatrics* 149, supplement 5, S1–S10. This synthesis proposes schools as platforms for coordinated education and health support; it does not test a HAO intervention. Several staged claim snippets are cross-column or reference-list artifacts. See the primary-PDF claim review (review retained privately). [icn-0048]
- **Beyond polarization towards dynamic balance: harmony as the core of mental health** — Antonella Delle Fave; Marié Philipina Wissing; Ingrid Brdar (13 September 2023). *Frontiers in Psychology* 14, 1177657. This “Hypothesis and Theory” article proposes a contextual and dynamic definition of mental health; it does not test HAO member outcomes. The staged claims contain cross-column fragments and a bibliography-only row; see the primary-PDF claim review (review retained privately). [icn-0054]
- **Near-optimal Integration of the Magnitude Information of Time and Numerosity** — Taku Otsuka; Yuko Yotsumoto (2023). *Royal Society Open Science* 10, article 230153. In two visual-judgement experiments with 42 volunteers, reliability-weighted integration approached a maximum-likelihood model more closely when time and dot-count information became available nearer together. Observed variability exceeded the model prediction in both experiments; the study does not establish optimal perception generally or test organisational decision-making. Bounded primary-PDF claim review (review retained privately). [icn-0052]
- **Gen Z job seekers should be ‘willing to do anything’, says Squarespace CMO Kinjil Mathur** — Orianna Rosa Royle (20 July 2024). This *Fortune* interview reports Mathur’s career advice and her account of unpaid internship work; it does not measure outcomes or describe HAO member practice. Bounded primary-PDF spot check (review retained privately). [icn-0094]
- **How FAST Feedback Boosts Performance** — Jeffrey Fermin (18 January 2024 in the archived version; the current page displays a later date). This AllVoices practice article proposes frequent, accurate, specific and timely feedback; its performance claims do not establish an evaluated FAST intervention. Bounded primary-PDF spot check (review retained privately). [icn-0061]
- **The Power of a Virtuous Cycle to Motivate Long-Term Goals** — Jean Moroney (2021; publisher page metadata). This practice article proposes a short-term payoff, a doable step and a value-oriented debrief for sustaining effort; it does not establish that this routine is necessary or effective for HAO members. Bounded primary-PDF spot check (review retained privately). [icn-0066]
- **Knowledge Brief: Basic Life Skills Curriculum** — UNICEF Azerbaijan (issuing context; n.d.). This three-page brief describes a curriculum developed with Azerbaijan’s Ministry of Youth and Sport for young people; its life-skills definition is attributed to WHO. The audit handoff assigns an unverified date and coauthorship; see the primary-PDF spot check (review retained privately). [icn-0117]
- **Building the Skills Adults Need for Life: A Guide for Practitioners** — Center on the Developing Child at Harvard University (n.d.). A four-page practitioner guide on planning, focus, self-control, awareness and flexibility. The handoff’s 2016 publication date is not established by the PDF; see the primary-PDF spot check (review retained privately). [icn-0107]

Legal Structures & Governance Instruments (8)

- **Choosing a Business Entity: A Guide for Worker Cooperatives** — Camille Kerr (2014), Democracy at Work Institute. This US worker-cooperative guide compares entity and tax choices, liability, governance and capital access. Its chart is a historical overview, not current jurisdiction-specific legal advice; a categorical double-taxation callout does not account for the LLC tax elections the guide itself describes. Bounded primary-source spot check (review retained privately). [icn-0017]
- **Thinking about starting a worker co-op? Here are some questions to consider before meeting with a lawyer** — Jenny Kassan (original-post URL dated 2010; PDF copy modified 2012). This one-page intake checklist poses 28 questions about jurisdiction, membership, governance, capital and tax concerns. It does not prescribe legal requirements or a HAO structure. Bounded claim review (review retained privately). [icn-0016]
- **Comparing Nonprofit and Cooperative Entities** — University of Wisconsin–Madison Center for Cooperatives (n.d.). This Wisconsin-oriented bulletin distinguishes state incorporation from federal tax status. Its absolute distribution language and suggestion of a cooperative tax “exemption” require qualification; see the primary-PDF claim review (review retained privately). [metadata incomplete] [icn-0040]
- **Essays on Dynamic Contracts** — Nan Zhao (2021). Boston University doctoral dissertation. Its three formal models address delayed evaluation, fixed-wage screening and dynamic delegation under distinct information assumptions. The generated claims overextend a Chapter 1 observability assumption to the whole dissertation and omit a condition on its full-effort result. The models do not test HAO contracts or governance. See the primary-PDF claim review (review retained privately). [icn-0045]

- **Application for Registration of INTEGRATED COOPERATIVE NETWORK** — PROVIDE.IO LLC (2024). This document is an Assumed Business Name application e-filed with the Oregon Secretary of State on 22 April 2024; it does not specify HAO ownership or governance. Bounded primary-PDF spot check (review retained privately). [icn-0002]
- **Contract Redlining** — DealHub Experts (updated 27 December 2023 in the archived copy). This vendor glossary describes collaborative edits and version control during contract negotiation. Its categorical description of redlining as a necessary step before signing is not a legal requirement; the publisher’s current FAQ says redlining is not always needed. [icn-0062]
- **Contract Versioning** — DealHub Experts; reviewed by Eyal Orgil (published 10 August 2023 in the archived copy). This vendor glossary describes tracking contract revisions, permissions and document versions. Its statements about enforceability and regulatory compliance are not independent legal authority. [icn-0063]
- **The Ultimate Guide to Company Structure Charts** — Mark Little (29 March 2023, per Lexchart’s author archive). This practical article distinguishes ownership, governance and management charts and proposes a five-step method for drawing them. Its examples are chart-design guidance, not evidence for a HAO governance arrangement. [icn-0064]

Technology, DLT & Networks (7)

- **SOFROP: Self-Organizing and Fair Routing Protocol for Wireless Networks with Mobile Sensors and Stationary Actors** — Mustafa İlhan Akbaş; Matthias R. Brust; Damla Turgut (2010). *Proceedings of the IEEE Conference on Local Computer Networks*, 456–463. The authors evaluate a routing protocol for mobile sensors and stationary actors through simulations of a particular wireless-network scenario. Bounded primary-PDF spot check (review retained privately). [icn-0074]
- **A Mobile Community Service Platform Promoting Ubiquitous Collaboration** — Kyungran Kang; Jungtae Lee; Kyoungwon Beak; Sungeun Park; Jaehwan Kim (2011). *Ninth IEEE International Conference on Dependable, Autonomic and Secure Computing*, pp. 939–946. The authors propose a phone-and-server architecture for role-based mobile communities and diagram example interaction sequences. The paper does not report a user study or measure collaboration outcomes; four of five staged evidence quotations splice its two columns. Bounded primary-PDF claim review (review retained privately). [icn-0075]
- **WebGIS based community services architecture by griddization managements and crowdsourcing services** — Haiyin Wang; Jianhua Wan; Zhe Zeng; Shengchuan Zhou (2016). *IOP Conference Series: Earth and Environmental Science*, 46, article 012048. The authors describe a WebGIS architecture for government community-service management in Huangdao District, Qingdao, with cloud-based service centres and mobile reporting. Bounded primary-PDF spot check (review retained privately). [icn-0079]
- **A Flexible Design for Funding Public Goods** — Vitalik Buterin; Zoë Hitzig; E. Glen Weyl (2019). *Management Science*, 65(11), 5171–5187. The canonical library PDF is the August 2020 arXiv revision, not the journal typesetting. It proposes quadratic finance under specified modelling assumptions; its deficit analysis qualifies the abstract’s optimality result. See the primary-PDF spot check (review retained privately). [icn-0103]
- **Cooperatives: An Ownership Model for Digital Networks** — Contributors affiliated with Orrick, Upside Cooperative and KPMG LLP (2023). The article proposes a cooperative ownership arrangement for digital networks and discusses conditional securities-law considerations; it is not a comprehensive compliance guide. Neither the staged “NSACoop” attribution nor the audit handoff’s a16z attribution appears in the PDF; see the primary-PDF spot check (review retained privately). [icn-0105]
- **Neurotechnologies: The New Frontier for International Governance** — Ricardo Chavarriaga; Jean-Marc Rickli; Federico Mantellassi (April 2023). *Geneva Centre for Security Policy, Strategic Security Analysis*, issue 29. This policy analysis surveys neurotechnology applications and governance proposals; it is not an OECD report or an empirical assessment of HAO. Primary-PDF spot check (review retained privately). [icn-0106]
- **The CrowdStrike Failure Was a Warning** — Brian Klaas (21 July 2024). *The Atlantic*, Ideas. This opinion essay uses the outage to argue for redundancy, diversity and slack in interconnected systems; it does not test those measures or HAO resilience. Bounded primary-PDF spot check (review retained privately). [icn-0095]

Other (Operating Models, Beyond Budgeting, Business Practice & Strategic Disclosure) (18)

- **Strategic Disclosure in the Patent System** — Douglas Lichtman; Scott Baker; Kate Kraus (2000). *Vanderbilt Law Review* 53(6), 2175–2218. The authors model when firms leading or trailing a patent race might publicly disclose research to affect a rival’s patent prospects, and consider whether private bargaining might displace disclosure. Its discussion of US first-to-invent priority and the former 35 USC § 102(b) describes the law at publication, not the current patent regime. [icn-0008]
- **The Fallacy of Premature Optimization** — Randall Hyde (2009). *ACM Ubiquity* 10(3), 24 February–2 March. A software-engineering opinion essay arguing that performance should be considered during system design while premature micro-optimisation is avoided. It offers no HAO-specific evidence. The essay’s attribution of the “97%” passage to Tony Hoare is not established by the cited article; Donald Knuth printed it in his own 1974 essay. Two staged rows quote an embedded passage by Charles Cook rather than Hyde’s own words. Bounded primary-PDF review (review retained privately). [icn-0011]
- **The Evolution of a Management Accounting Idea: The Case of Beyond Budgeting** — Sebastian Becker; Martin Messner; Utz Schäffer (19 February 2009 manuscript; repository wrapper labels January 2010). A qualitative genealogy based on publications, interviews and observation examines the tension between Beyond Budgeting as a complete management model and adaptation of its component practices. The authors offer one interpretive explanation for its limited diffusion, not a causal or HAO-specific test. The generated summary misstates their conclusion; see the full primary-PDF claim review (review retained privately). [icn-0072]
- **Operating Models: Delivering on Strategy and Optimizing Processes** — Annie Murphy; Jamie Kirwin; Khalid Abdul Razak (© 2016 Ernst & Young LLP). This EY practitioner guide frames an operating model as a bridge between strategy and operations and proposes adaptable design elements. Its performance benefits are argued rather than measured, and the staged claim set contains a duplicate-ID collision; see the primary-PDF claim review (review retained privately). [icn-0032]

- **Introducing the Next-Generation Operating Model** — Digital McKinsey (PDF colophon: January 2017). This is an article collection. The staged claims mostly come from “The next-generation operating model for the digital world” (printed pp. 8–15), whose canonical-edition byline names Albert Bollard, Sanjay Kaniyar, Elixabete Larrea, Alex Singla and Rohit Sood; one claim comes from the introduction. The prior Bucy–Hall–Yakola attribution does not match this PDF. The collection offers practitioner arguments and anonymized cases, not a test of HAO; see the primary-PDF claim review (review retained privately). [icn-0033]
- **Beyond Budgeting** — Vassili Ioannidès de Lautour (2018). In *Strategic Management Accounting, Volume I: Aligning Strategy, Operations and Finance* (pp. 217–270). Palgrave Macmillan. The chapter analyses criticisms of master budgets and describes Beyond Budgeting as an alternative management philosophy. The ten generated claims mostly reproduce the author’s conceptual critiques, not measured effects; the four-month budget-preparation figure is a rounded secondary estimate. See the full primary-PDF claim review (review retained privately) and earlier numerical spot check (review retained privately). [icn-0082]
- **Understanding Why Beyond Budgeting Has Not Been Widely Adopted** — Phil Hudson (2012). *Doctor of Business Administration thesis, Kingston University*. An exploratory survey of 185 management accountants and 50 follow-up interviews found partial uptake of component practices among participants while fixed budgets remained common. The author cautions that the sample cannot support generalisation. The canonical PDF’s 2019-09 filename does not give the thesis date. See the full primary-PDF claim review (review retained privately). [icn-0102]
- **An Empirical Investigation of Beyond Budgeting Practices** — Michal Matějka; Kenneth A. Merchant; Winnie O’Grady (April 2020 preprint; journal version 2021). An exploratory survey compares 80 self-reported partial or full adopters with 121 non-adopters drawn from Beyond Budgeting Round Table contacts. The convenience sample and self-reported measures limit generalization; observed associations do not establish causation. The canonical PDF’s 2022-05 filename does not give the paper date. See the full primary-PDF claim review (review retained privately). [icn-0085]
- **An Exploratory Study on Corporate Governance From Neuro-Governance Lenses in the Malaysian Context** — Larisa Ivascu; Codruta Daniela Pavel; Muddassar Sarfraz; Benedict Valentine Arulanandam; Hong Yip Tan (2022). *Frontiers in Psychology*, 13, 911907. A qualitative study reporting 21 structured interviews about corporate misconduct and related decision-making; its literature-review examples and interview opinions do not establish general human-behaviour laws. The audit handoff misidentified the authors; see the primary-PDF spot check (review retained privately). [icn-0104]
- **Beyond budgeting in practice** — Agnes Burell; Anna Mattsson (2022). *Master’s thesis, Department of Business Studies, Uppsala University*. A four-case qualitative study based on eight semi-structured interviews. Although its abstract calls the cases retail organisations, one is a pharmaceutical-company pilot. The staged claims[] extraction covers the abstract and background/theory sections, not the empirical findings; see the full primary-PDF claim review (review retained privately). [icn-0086]
- **How Operating Models Must Evolve for the Future** — Steve Basili (EY; 5 April 2022). EY-sponsored MIT Sloan Management Review article on proposed operating-model changes for consumer-products companies. It draws on a separately reported survey but does not itself establish that these recommendations improve performance. The canonical PDF is a printout of the sponsored web page, not the survey report; see the primary-source spot check (review retained privately). [icn-0060]
- **Iterative Business Model Innovation: A Conceptual Process Model and Tools for Incumbents** — Nadine Bachmann; Herbert Jodlbauer (2023). *Journal of Business Research*, 168, 114177. A structured narrative review of 47 publications develops a six-phase conceptual process model with 23 activities and 38 tools; it does not test HAO outcomes. Five staged claim rows are reference-list or author-biography artifacts; see the primary-PDF spot check (review retained privately). [icn-0088]
- **Strategic Disclosure and Investor Loss Aversion** — Zeqiong Huang; Joseph D. Piotroski; Chloe L. Xie (November 2023 working paper). A stylised public-company model predicts disclosure of mildly bad earnings news and withholding of mildly good news under specified investor-loss-aversion conditions. It does not test HAO governance. The staged source record bears an unrelated title and agent-based summary; see the full primary-PDF claim review (review retained privately). [icn-0055]
- **Guide to the Systems Engineering Body of Knowledge (SEBoK), version 2.10** — SEBoK Editorial Board; Nicole Hutchison, editor in chief (released 6 May 2024). This PDF extraction of the SEBoK wiki covers systems-engineering foundations, management, applications and related disciplines; its articles have separate contributors. The version and editorial responsibility are documented in the primary-PDF spot check (review retained privately). [icn-0093]
- **Integration vs Implementation: The Value of Each** — DecisivEdge (site byline “The Story Dev”; 2 October 2023). This company blog article distinguishes IT implementation from systems integration. Its claim that 70% of CRM implementations fail names no supporting study and is not used as evidence here. [icn-0090]
- **(Pro-)Social Learning and Strategic Disclosure** — Roland Bénabou and Nikhil Vellodi (14 May 2024 draft). In their sequential-experimentation model, disclosure is polarised and positively selected when prior uncertainty is large; that conditional result is not an observed HAO outcome. A revised journal article appeared in 2025. Bounded primary-PDF spot check (review retained privately). [icn-0091]
- **Strategic Disclosure with Reputational Concerns** — Wenhao Zhang (2024). *Journal of Mathematical Economics* 111, 102945. The author models an expert’s disclosure incentives when honesty and reputation matter. The canonical PDF is a publisher-page printout with the abstract and introduction but only snippets of the model, analysis and conclusion; it is not the full article. Its theoretical result is not evidence of HAO outcomes. See the bounded primary-source claim review (review retained privately). [icn-0057]
- **What is an Operating Model?** — Quantive (publisher; no named author or publication date; archived 18 April 2024). This company article distinguishes operational design from a business model’s commercial activities and sets out six operating-model components. Its account of a “Modern Operating Model” describes Quantive’s proposed approach, not independent evidence of its results. [metadata incomplete] [icn-0065]

5. Excluded

Two documents in the source PDF collection were excluded from the research library above as non-bibliographic:

- icn-0096 — a CliffsNotes study guide for *Lord of the Flies*. Confirmed by its own machine-extracted title (“Cliffsnotes Lord Of The Flies”) and summary (literary analysis of the novel). Unrelated to the corpus’s subject matter. Named here so a later pass does not re-add it.

- `icn-0116` — a **Docusaurus site-configuration file for the provide.io governance website**. A PDF snapshot of a `docusaurus.config` file (site URL, organization name, docs and blog paths), confirmed by its machine-extracted title (“provide.io governance”) and summary (site-configuration details). A build artifact, not a reference work. Named here so a later pass does not re-add it.

A full filename scan for junk-pattern keywords (cliffsnotes, sparknotes, wikipedia, test, sample, lorem, placeholder) across all 118 canonical documents found no other candidates; these two are believed to be the complete set within this collection.

The source collection’s index was subsequently repaired. `icn-0118` now has a row in `indexes/library-index.csv`, including its canonical-file checksum and source listing; the index contains 118 data rows. This document remains one of the 116 works included above, because two non-bibliographic files are excluded. The prior absence of this row was a catalog gap, not a reason to omit the work.

Research Evidence Map

Purpose and authority

The 118-source/929-claim extraction inventory is a discovery and review aid. The extraction inventory originates from the private research library. Generated claims, summaries, labels, and selected quotations are discovery aids only. The inventory preserves generated extraction output so that later review can locate material, but it does not establish the accuracy, relevance, adoption, or authority of any extracted item.

Extraction QA acceptance is not corpus authority. Direct verification against the canonical primary document remains required before corpus promotion. A reviewer must check the relevant material in context before it can support a chapter claim, an appendix conclusion, or a lineage status. No lineage entry or chapter claim has been promoted in this phase.

Coverage

The preservation inventory has 118 sources; the bibliography represents 116 works; two preservation records are intentionally excluded from the bibliography. The inventory also contains 929 claims and 718 evidence snippets. Extraction QA records 107 QA-accepted extraction records and 11 low-evidence records; the associated claim split is 896 accepted and 33 low-evidence claims.

The promotion-status split is 105 discovery-only, 11 quarantined-low-evidence, and 2 quarantined-excluded source records. For claims, it is 881 discovery-only, 33 quarantined-low-evidence, and 15 quarantined-excluded claims. Its claim disposition is 914 represented claims and 15 intentionally excluded claims.

`icn-0096` and `icn-0116` remain excluded from the bibliography and the review queue. They remain in preservation material so that the inventory is auditable, but neither is a candidate for corpus use.

Relationship to the corpus

The bibliography in `references.md` identifies the represented works. `lineage.md` records the current ancestry statuses, and `source-integration-register.md` states the source-authority boundary. `open-questions.md` remains the register for unresolved corpus positions. This inventory changes none of those positions.

First primary-document review queue

Cluster	Source IDs queued	Review question
Polycentric governance and Ostrom	<code>icn-0013</code> , <code>icn-0030</code> , <code>icn-0035</code> , <code>icn-0042</code> , <code>icn-0058</code> , <code>icn-0081</code> , <code>icn-0087</code>	Which material survives contextual primary-document review, and which identifies conceptual correspondence without ancestry? The extraction set contains no work titled or authored by Ostrom; direct Ostrom verification must use the canonical bibliography, not this ledger.
	<code>icn-0071</code>	Does the source support a precise comparison with the HAO structure without implying historical influence?
Beer’s Viable System Model	<code>icn-0068</code> , <code>icn-0069</code> , <code>icn-0077</code> , <code>icn-0083</code>	Which autonomy, cooperation, and recursion properties correspond to HAO units?
Holonic manufacturing	<code>icn-0072</code> , <code>icn-0082</code> , <code>icn-0085</code> , <code>icn-0086</code> , <code>icn-0102</code>	Which devolved-authority material is supported, and what limits or adoption failures qualify it?
Beyond Budgeting	<code>icn-0016</code> , <code>icn-0017</code> , <code>icn-0037</code> , <code>icn-0040</code>	Which entity-form statements remain accurate in their stated jurisdiction and publication context?
Worker-cooperative law		

Review rule

Queue membership sets review priority only. It does not establish verification, adoption, truth, support for a corpus position, or evidence of ancestry. A claim that fails direct primary-document verification is rejected rather than weakened into an unattributed corpus assertion.

HAO Source Integration Register

A record of HAO-related material held outside the canonical corpus, its authority, and its disposition.

1. Scope and path convention

This register prevents auxiliary, generated, and implementation-dependent material from acquiring authority merely because a filesystem search finds it. The canonical HAO corpus remains the `docs/` tree in the HAO repository. Each source below is classified independently, and integration means terminology or metadata synchronization unless the entry states that claim-level review occurred.

Repository sources use compact visible labels so the register remains readable in narrow and printed layouts. Public links target the exact reviewed commit. Every row also provides a native “Full source identity” disclosure containing its complete `repository@<full-commit>:<repository-relative-path>`, local-repository, workspace, or package identity; the print edition expands that metadata without replacing the compact label. The public site renders private repositories and local-only holdings as provenance text rather than links that an unauthenticated reader cannot open. Links in the canonical-destination column resolve from this appendix in the canonical repository and are translated to public routes, or identified as local-only, when the site imports the register. The HAO archive entries are reviewed at `d3b8581b39349ffc41683375f853071b129f8031`, the source revision containing the migrated archive paths; later Task 5 register commits are not represented as external source revisions.

2. Source register

Source path	Authority	Content class	Canonical destination	Integration status	Last reviewed	Notes
that-explains-it@b9afe399.../manifesto	Derived copy; non-authoritative	HAO explanatory prose	Consolidated manuscript and canonical economics chapter	Terminology synchronized; retained as an auxiliary example	2026-08-23	Configured origin: <code>git@github.com:livingstaccato/that-explains-it.git</code> . Directionality distinguishes outward allocation from diminishing inward contributions. Content changes do not flow back into the corpus without claim-level review.
that-explains-it@b9afe399.../manuscript	Derived copy; non-authoritative	Consolidated-manuscript example	Consolidated manuscript	Terminology and directionality synchronized with the canonical manuscript; retained as an auxiliary example	2026-08-23	The HAO repository remains authoritative when the copies diverge.
that-explains-it@b9afe399.../captured research PDF	Artifact-original captured research; no claim authority	Captured research artifact	Canonical local PDF archive (local archive; not publicly linked)	Metadata filename normalized; helper copy retained	2026-08-23	The rename did not rewrite the binary: 907,470 bytes; SHA-256 <code>e25f7894335a135ac19ae83deb93001c3b6ed6756bbbfefeeccdaed9aaf41bef4</code> before and after. Its claims and citations require source-level verification before use.
provide-scrap@3ddb1d2b.../research report	Generated research report; no claim authority	Research input requiring claim-level review	Canonical economics chapter	Terminology and directionality synchronized; provenance advisory added; registered as background only	2026-08-23	Local repository with no remote configured. Source URLs, citations, figures, and claims have not received claim-level verification. No claim or figure was promoted into canonical prose, and the report is not primary evidence.
workspace/icn-text-hits-2026-05-31.txt	No evidentiary authority	Dated legacy-search inventory	This register; source claims must resolve to their canonical files before use	Preserved in place as a historical inventory	2026-08-23	Non-repository holding with no reproducible generator or source list found in adjacent scripts, READMEs, filename references, or available shell histories. The neutral filename and original bytes were retained intentionally: 30,888,302 bytes; SHA-256 <code>6ee782c659c53ec8c68e871546affa4db0243c65d3a9e307ffc09ea967f22411</code> . Its embedded excerpts and absolute paths describe the 2026-05-31 search state, not integrated content.
hao-scrap@ d3b8581b.../research PDFs (private repository; not publicly linked)	Original binary evidence; not canonical prose	Imported PDFs in the local research archive	NotebookLM local archive (local archive; not publicly linked)	Registered; filename and metadata normalized; integrity metadata incomplete	2026-08-23	Configured origin: <code>git@github.com:livingstaccato/hao-scrap.git</code> . Pre- and post-migration hashes verified byte identity for files whose hashes were captured during the migration. Archive manifests do not consistently persist PDF checksum and byte-count fields: the matching NotebookLM source records are metadata-only, and the two media PDFs have no manifest hashes.

Source path	Authority	Content class	Canonical destination	Integration status	Last reviewed	Notes
hao-scrap@d3b8581b.../audio (private repository; not publicly linked)	Original binary evidence; not canonical prose	Imported audio and lossless working copies	Archived audio (local archive; not publicly linked) and lossless working archive (local archive; not publicly linked)	Archived locally with normalized path and metadata names	2026-08-23	Recorded speech remains the evidentiary original. Filename and metadata normalization does not alter audio bytes; transcripts carry the separate editorial notice required by the archive policy.
protobuf@74211c0d.../SingleFieldBuilder.java	External project authority; outside HAO ownership	Vendored protobuf implementation dependency	None; excluded from the HAO corpus	Verified exclusion; not edited	2026-08-23	The same unrelated implementation-comment match was reviewed in two local protobuf checkouts, both at this detached revision with the configured origin https://github.com/protocolbuffers/protobuf.git .
torch 2.13.0/schedules.py	Package-distributor authority; outside HAO ownership	Installed Python dependency	None; excluded from the HAO corpus	Verified exclusion; not edited	2026-08-23	Stable distribution identity is <code>torch==2.13.0</code> ; installed package metadata names https://github.com/pytorch/pytorch as its repository. One unrelated pipeline-scheduling comment matches the separator-wildcard audit. Virtual-environment content is not owned source.

Full source identities

- `livingstaccato/that-explains-it@b9afe399ea82153bb32ec1f69cf3d215062335b9:examples/hao-a-multi-level-meme-manifesto-works.md`
- `livingstaccato/that-explains-it@b9afe399ea82153bb32ec1f69cf3d215062335b9:examples/humanized-autonomous-organization.md`
- `livingstaccato/that-explains-it@b9afe399ea82153bb32ec1f69cf3d215062335b9:examples/chatgpt-2025-04-trickle-out-economics-and-bottom-up-models-historical-examples-and-modern-applications.pdf`
- `local-repository:provide-scrap@3ddb1d2b83a71333a2b7742ec867c496180babab:20250321-1412-gemini-deep-research-report.md`
- `workspace:code/gh/provide-io/icn-text-hits-2026-05-31.txt`
- `livingstaccato/hao-scrap@d3b8581b39349ffc41683375f853071b129f8031:notebooklm-livingstaccato/{chatgpt - 2025-04 - Trickle-Out Economics and Bottom-Up Models_ Historical Examples and Modern Applications.pdf,media/*.pdf}`
- `livingstaccato/hao-scrap@d3b8581b39349ffc41683375f853071b129f8031:notebooklm-livingstaccato/{media/{Flipping_the_pyramid_with_trickle_out_economics,The_Blueprint_for_Trickle_Out_Investment,Venture_Capital_versus_Trickle_Out_Capital}.m4a,_work/wav/*.wav}`
- `protocolbuffers/protobuf@74211c0dfe2777318ab53c2cd2c317a2ef9812de:java/core/src/main/java/com/google/protobuf/SingleFieldBuilder.java`
- `torch==2.13.0:torch/distributed/pipelining/schedules.py`

3. Authority boundary

Auxiliary prose can identify material for review, but it cannot establish a corpus claim. Generated research requires source-by-source verification, captured binaries preserve what a research system received or produced, and search inventories record discovery state only. A claim enters the HAO corpus only through review against its underlying source and an explicit change to the relevant canonical chapter or appendix.

HAO Framework Terminology (Extended Glossary)

An authoritative reference for all core concepts, entities, and constructs in the Humanized Autonomous Organization model.

I. Structural Entities

Term	Definition	Plain name
HAO (Humanized Autonomous Organization)	A socio-technical framework that centers human agency, polycentric governance, and ethical alignment while integrating AI, economic coordination, and trust systems.	the network’s coordinating framework
ICN (Integrated Cooperative Network)	A specific instantiation of the HAO model designed for trickle-out economics and distributed entrepreneurship via UMEs and SEPs.	the reference cooperative business network
MTU (Member Trust Union)	A cooperative financial infrastructure layer of the ICN, designed to reimagine banking around trust-centered services, decentralized credit, and financial sovereignty.	the network’s credit-union-like financial institution
UME (United Micro Enterprise)	The smallest operational unit in the HAO/ICN, typically a semi-autonomous, self-governing team or business with embedded ethical, financial, and governance protocols.	a small, self-managing venture team (≤ ~15 people)
SEP (Strategic Enterprise Partnership)	A time-bound or goal-specific alliance between two or more UMEs formed to pursue a shared opportunity, project, or market interface.	a joint venture between teams
MEE (Micro Enterprise Ecosystem)	The protected internal economy in which UMEs operate, distinct from (but connected to) external markets through defined interfaces.	the network’s protected internal economy

II. Governance and Agreements

Term	Definition	Plain name
AGF (Adaptive Governance Framework)	The multi-level, polycentric governance model in HAOs that enables distributed authority while maintaining value coherence.	the network's layered governance system
DEA (Dynamic Enterprise Agreement)	A versioned, evolving governance charter that replaces fixed bylaws; it formalizes the rules, roles, and operational principles of a HAO or UME.	a versioned operating agreement replacing fixed bylaws
IR (Intentional Redundancy)	The design principle of overlapping roles, responsibilities, or systems to increase resilience and fault tolerance without sacrificing efficiency.	
EIA (Enterprise Integration Assessment)	A formalized process for evaluating whether and how external entities (individuals, UMEs, organizations) can be integrated into the ICN or HAO.	
VAM (Value Alignment Monitoring)	Systems or protocols designed to continuously track whether behaviors, outputs, and decisions align with the stated principles of the HAO/ICN.	ongoing checks that actions match stated principles
Governance Node	A locus of decision-making authority in a polycentric network, whether at the UME, regional, or network-wide level.	
Consent Governance	A decision-making model focused on resolving objections rather than achieving majority rule; often used for proposals and conflict mitigation.	

III. Financial Architecture

Term	Definition	Plain name
Trickle-Out Economics	A financial design where capital flows from the coordinating entity (HAO) outward to the edge (UMEs) based on value-creation potential. A separate diminishing contribution protocol returns a declining share of revenue to shared infrastructure as UMEs mature.	
Mutual Credit System	A network-internal financial mechanism in which credits are created at the moment of value creation and circulate without requiring fiat injection.	
Revenue Allocation Framework	A structured, milestone-based protocol defining how revenues are shared between UMEs, SEPs, and the HAO over time.	
SEP Revenue Contract	A dynamic revenue-sharing agreement specific to a SEP that defines contribution, risk, trust weight, and revenue distribution logic.	
Distributed Secure Credit	A decentralized form of credit issuance where trust, not collateral, underwrites lending, and risk is spread across the network.	
Resilience Reserve	A portion of revenue set aside by the HAO or UMEs to support recovery, failure mitigation, or reinvestment in emergent opportunities.	
Public Market Interface	A buffer mechanism (like Contribulo) that allows regulated interaction between the MEE and traditional financial markets, insulating UMEs from external volatility.	a buffer company between the network and outside investors

IV. Trust, Identity, and Participation

Term	Definition
Trust Graph	A weighted, dynamic graph structure representing the multi-dimensional trust relationships between agents (members, UMEs, SEPs) in the HAO.
Progressive Trust Verification	A non-binary, layered approach to trust wherein permissions, roles, and credit access grow as reputational signals accrue over time.
Self-Sovereign Identity (SSI)	A decentralized identity system where individuals control their credentials, proofs, and disclosures without reliance on central authorities.
ZKP (Zero-Knowledge Proof)	A cryptographic method allowing users to prove something (e.g., experience level) without revealing personal details, used in trust-preserving verifications.
Trust Horizon	The maximum range of transactional or governance influence an agent can have based on their trust profile and relationship map.
Reputation Inertia	The tendency of reputation systems to resist updates after infrequent or outdated participation; countered by time-decay mechanisms in HAOs.
Trust Decay Curve	A rate at which unused or unrefreshed trust scores lose weight, ensuring that reputation remains current and context-relevant.

V. Technology and Intelligence Systems

Term	Definition
CIN (Collaborative Intelligence Network)	A network layer of AI agents, data systems, and human feedback loops that augment decision-making, value monitoring, and governance workflows.
AI-Augmented Governance	The practice of using AI agents as advisors, validators, or summarizers within governance—without removing final human authority.
Ethical ML Pipeline	A monitored machine learning architecture in HAOs that is trained on value-aligned data and outputs interpretable decisions for human governance.
Collective Memory System	A structured digital memory of proposals, outcomes, and decisions that informs future deliberations and captures organizational evolution.
Deliberation Agent	An LLM or rules-based bot tasked with surfacing overlooked perspectives, summarizing positions, and reducing cognitive load in decision-making.
Dynamic Risk Model	A live-adjusting risk framework that recalibrates credit or governance exposure based on behavior, context, and network-wide metrics.

VI. Lifecycle and Health Management

Term	Definition
UME Lifecycle Management	A process framework for tracking and supporting a UME from inception to maturity, potential degradation, or exit.
Enterprise Degradation	The process by which a UME or SEP drifts from HAO-aligned principles due to cultural, financial, or operational misalignment.
Enterprise Collapse	A terminal state where a UME ceases operation due to internal failure, misalignment, or systemic risk exposure.

Term	Definition
Antifragility Event	A network-level stressor that, rather than destroying the system, triggers learning, reconfiguration, and resilience strengthening.
Intentional Cannibalization	A design pattern where the HAO proactively dissolves or restructures outdated UMEs to seed new, better-aligned entities.
Health Signal Aggregator	A CIN node or tool that collects multi-modal data on cultural, financial, operational, and trust metrics to assess HAO health.

VII. Meta-Frameworks and Principles

Term	Definition
PARTS Model	The five-part foundational architecture of the HAO: Participatory Evolution Model (PEM), Adaptive Resilience Model (ARM), Resilience Operations Model (ROM), Transformative Infrastructure Model (TIM), and Systems Integration Model (SIM).
ETHICAL Framework	A values foundation for all HAO design: Empathy, Transparency, Harmony, Integrity, Collaboration, Accountability, Learning & Longevity.
MMM (Members, Mission, Market)	A priority ordering for what a UME attends to: members precede mission, because the people present outlast any particular mission; mission precedes market, because market opportunity alone does not determine what is worth building; market is the viability test the first two must pass. The framework's source documents are marked as working documents, not yet reviewed, and its content is treated as provisional here; see <code>docs/appendices/open-questions.md</code> for the framework's own open questions.
HAO Ontology Layer	A shared semantic layer describing roles, relationships, functions, and transitions within the HAO system.
Participatory Evolution Model (PEM)	A model that emphasizes continuous refinement of governance through structured proposals, reflection, and consensus loops.
Adaptive Resilience Model (ARM)	A model for absorbing disruption and reconfiguring afterward, directed at shocks that cannot be enumerated in advance rather than at risks that can: proactive stress testing, scenario modeling across disruption types, crisis-response training, and redundancy in key systems. Its source frames it as deliberate exposure to measured stress to surface vulnerabilities that abstract analysis misses, and treats some inefficiency as a hedge. \$VI's Antifragility Event is the same concept under another name. \$VI's Intentional Cannibalization is related but not identical: it dissolves units judged outdated, where this source proposes dismantling parts of a thriving unit during abundance, to keep attachment to any single asset in check. ARM and ROM are not cleanly distinguished in the source material — the archive's ROM brief reproduces most of the ARM brief verbatim, still naming ARM throughout — and this glossary does not supply a distinction the sources do not make. Imported from <code>provide.coop/docs/parts/02-adaptive-resilience-model.mdx</code> , an unreviewed brainstorming brief; see \$VII note below.
Resilience Operations Model (ROM)	A model for resource allocation and decision-making under uncertainty or constraint—favoring redundancy, adaptability, and fairness.
Transformative Infrastructure Model (TIM)	A model governing how the network builds and procures technology so that the technology carries the network's stated commitments rather than eroding them. Its stated scope is wider than user-facing software: algorithm design, data ownership, and vendor and hardware supply chains. Its source proposes screening vendors against a values standard, auditing the network's own tools for effects that are technically efficient but reduce member agency, accepting deliberate inefficiency or reduced data collection where efficiency would cost more than it returns, and monitoring the network's own systems for unintended consequences, the last of these framed by the source as watching for the pattern \$9.1.3 catalogs as Mission Creep. Imported from <code>provide.coop/docs/parts/04-transformative-infrastructure-model.mdx</code> , an unreviewed brainstorming brief; see \$VII note below.
Systems Integration Model (SIM)	An architectural model that monitors and manages interfaces between the HAO and external systems (e.g., markets, governments, protocols).

Note on the ARM and TIM entries. The corpus names all five PARTS models but describes only three of them in its own text. ARM and TIM appear in `~/code/hao/docs` only inside the PARTS Model row above, so their entries were imported from the source archive at `~/code/site-provide-coop/provide.coop/docs/parts/`.

Those two briefs are AI-generated ideation, not ratified specification. They are written largely as open questions (“Could the ICN...”), they use the earlier ICN vocabulary rather than the current terms, and several of their proposals are posed rather than adopted. The entries above record what each model is for and the mechanisms its brief names; they do not adopt those proposals as network policy. Both entries are narrower than their sources, and should be read as the least the source supports rather than the most. What is left out is chiefly the ARM brief's stress-test scenarios and the TIM brief's “slow tech” brand positioning; the TIM brief's vendor-screening mechanism is kept above, minus its “blacklist” framing.

This does not make ARM and TIM weaker entries than their neighbors. PEM, ROM, and SIM are no better attested: ROM appears nowhere in the corpus outside this table, and PEM and SIM appear only in a single chapter 09 passage that glosses them from this table. All five rows of this section trace to material no one has verified. The difference is that ARM and TIM name the file they came from, and it can be read; the other three do not.

VIII. Common Failure Modes (and Mitigations)

Term	Definition
Alignment Drift	A slow, unintentional shift away from founding principles; mitigated by VAM and CIN interventions.
Governance Capture	A scenario in which a powerful subset controls decision-making; addressed via trust decay, consent models, and rotating facilitation.
SEP Entropy	The disorganization or confusion that arises in multi-UME projects without clear agreements; solved through enforceable DEA branches and structured retrospectives.
Credit Inflation	A collapse in mutual credit value due to unchecked issuance or low circulation; mitigated with issuance limits, time-bound credits, or decay mechanisms.
Reputation Exploits	Gaming of the trust system by collusion or sybil-like tactics; requires graph analysis and anomaly detection tooling.