Distributed Inference

Introduction

Across history, power structures periodically destabilise and reconfigure as their coordination costs, legitimacy failures, or incentive models break down. Emerging solutions require a fundamental rethink of authority, merit, and value structures. Communities that embody these changes can be described as post-democratic or participatory economies, as they function through interdependent, trust-based, and meritocratic governance, without the need for arbiters or top-down authority.

In these communities, day-to-day life is the governance model itself; their operations and values continuously evolve through active participation. Historically, such communities have been rare because they require a dynamic balance between autonomy and shared values. They must adapt fluidly, whilst ensuring no single individual exploits collective autonomy.

Imagine community values were sealed in an envelope, and the governance system was tasked with continually deciphering and adapting to those values, serving the community's ever-evolving priorities without requiring constant explicit instruction.

The underlying tenet of Distributed Inference is: people who love their community and environment, and are represented in their community, contribute to their community, and therefore, in this context, to the wider network.

Social identity theory (Tajfel & Turner) holds that identification with a group motivates prosocial behaviour toward it. Robert Putnam's work on social capital shows that community embeddedness correlates strongly with civic contribution. Self-determination theory suggests that feeling represented and valued increases intrinsic motivation to participate.

Intentional communities are essentially natural experiments in this principle. Kibbutzim, co-housing movements, religious communities (Amish, monastic orders). Ecovillages consistently show high contribution rates, and the common factor researchers identify is the combination of genuine belonging and meaningful representation in decision-making. People build, maintain, sacrifice, and create for communities they feel are genuinely theirs.

The principles underlying DI have been observed working in practice across centuries and cultures by researchers studying commons that survived. Elinor Ostrom's Nobel Prize-winning analysis of commons governance identified the same design principles this document outlines. That convergence is offered as the most honest available answer to the question of whether any of this has ever actually worked. A generalisation of Ostrom asks you to explain why centuries of documented working commons across radically different cultures would suddenly stop working when the same principles are formalised.

Scale and Stakes

DI operates at multiple scales. At its smallest, a single DIM (Distributed Inference Model) might govern a housing cooperative of 20 people. At larger scales, DIN (Distributed Inference Network) becomes a mesh of interconnected DIMs: geographic communities and functional services coordinating laterally without centralised control. This could span thousands of communities and

potentially millions of people through gradual adoption as communities choose systems that serve them better than existing alternatives.

We are at a critical juncture. Representative democracy designed for 18th century constraints cannot handle the coordination challenges of unprecedented complexity: climate requiring global cooperation, economic concentration, institutional trust collapse, and information warfare. Too often, the responses: authoritarian strongmen, revolutionary collectivisation, and mass apathy are historically catastrophic. DIN provides a another option: mutual, adaptable self-coordination without centralised control, participatory governance that can scale.

The Problems with Traditional Systems

Representative politics functions by creating a mandate, which is a static representation of "what is good for us" frozen at the moment of an election. This mandate is a logic imposed from the top down, often ignoring the shifting reality of the people it claims to serve.

Traditional parties operate on prescriptive logic. They decide what the public needs and drive implementation. Whether it's a colonial administrator in the 20th century, an NGO in Sub-Saharan Africa, or a Westminster politician today, the impulse is: I have a superior logic, and I will impose it for your own good. When you impose a "Truth" from above, you aren't actually changing the being of the people on the ground; you're just forcing their being into a subterranean state. The logic stays on the surface, while the actual state of the people is not accurately honoured.

The assumption that you can pre-know the answer to fair governance is not just intellectually arrogant; it is self-defeating, because the act of pre-knowing forecloses the sampling process through which you might discover whether your answer was correct. Representative democracy commits this fallacy, as does a founding ideology that defends against the pressures of reality.

Society is currently captured into a closed loop because it is built on fear and accumulation. Monopolies and their clauses are an attempt to hold against loss, to instil an immovable status quo with the pressure of their self-bias. The commons from which the skills and generative knowledge are sourced is seen as a surface of capture for the sake of capital protection.

Representative systems did not become capital-first by accident. The funding of elections, ownership of media ecosystems, and the architecture of lobbying were not corruptions of an otherwise neutral system, but acts of self-interest by people with the resources and foresight to shape the rules of the game before the game was settled. Self-interest operating at the level of institutional design is blind to itself because it is itself. It looks like adaptable sensemaking. The entrenchment of power shapes norms in society, instilling the shape of the frame the next generation graduates into, pre-baking the meta of what is rational to suit the preexisting shape—under the actually unproven assumption of its merit—and with its merit seemingly proven only by its apparent success—a perfect circle. It transforms from manipulation into ambient assumption, until the interests of those who built the frame become indistinguishable from the frame itself, and challenging them requires challenging something that no longer presents itself as a distinct position. Narrow self-interest does not merely shape institutions passively. It actively degrades the conditions that might produce alternatives, weakening the commons, atomising communities, and defunding the civic infrastructure through which collective judgment might otherwise form. A population that cannot coordinate cannot exit, and a population that cannot exit cannot negotiate. Dependency is not an accident of the system. It is one of its primary products. Selecting over time the kinds of people and proposals that align with its sensemaking norms, filtering out those that do not, before any

demand for coercion arises.

The Overton window is as epistemological as it is political. The range of possibilities that concern us as inhabitable rather than merely conceivable is siloed to the point where genuine alternatives read as naivety or danger. The pre-emptive embarrassment of alternatives crushes experimentation.

Ivan Illich's argument was that beyond a certain threshold, institutional provision of human needs does not supplement human capacity but displaces and eventually destroys it, merely by substitution.

The Problem of Structurelessness

Some communities attempt governance through informal norms only: "no rules, just freedom." What often happens: a charismatic person emerges naturally and accumulates informal authority through social dynamics. They become de facto leader with no formal accountability because "no rules." The result is either dictatorship or capture. The community stays small and homogeneous, collapses from internal conflict, or gets captured by external forces.

Structurelessness is a structure, one that's invisible, unaccountable, and favours certain personality types: the socially dominant, persistent, or those who think strategically. In the absence of explicit governance, implicit governance emerges and concentrates power in ways that can't be challenged or even clearly identified.

Even perfectly functioning internal governance must interface with external systems: taxes, property, contracts, suppliers, services. Legal entities, bank accounts, authorised signatories, stable negotiating parties are required, which itself becomes a power concentration point.

If you start with the premise that subjectivity is an infinite, intractable wall, then any coordination is impossible. But the fact that we can even have a followable conversation proves that there is a common ground. The meta of all subjective logic is infinite, but the overlap humans exist in is significant enough to have built a shared society.

Freedom as naively thought of cannot exist for any individual of a collection that share finite resources and space, each "freedom" is a negotiated dependency, constantly effected by the churn of environmental consequences created by other individual "freedoms." There is no space in which your freedom is not already someone else's limitation, in some fashion, no matter how seemingly distant.

The Problems with Pure Consensus

Pure consensus has critical issues at scale without injecting the very hierarchy or information structures it aims to circumvent. It bogs down, gets captured by the most patient or manipulative voices, or dissolves into process worship. Decision latency so tepid that life moves on before time.

The biggest issue with pure consensus models is unaccountability: who enforces decisions when nobody has defined authority?

Decision Paralysis: When everyone has equal power, situations can arise where nobody feels responsible for enforcing decisions or culling discussion. Consensus cost grows superlinearly with group size.

Lack of Formal Mechanisms for Dispute Resolution: Without clear authority structures, social conflicts remain unresolved. Members hesitate to challenge issues for fear of causing conflict. A member subtly exploits others, but nobody speaks up because it's unclear who should intervene.

Authority Cascade: In order to rise to an attempt of authority by another, one has to posit an opposite position with power, which can then be objected to by another.

Selection Bias & Insular Thinking: Communities tend to attract like-minded individuals, unintentionally filtering out diversity of thought. This creates intellectual blind spots that limit innovation. A group promoting "holistic health" may reject valuable scientific insights because they unconsciously select members who reinforce existing beliefs.

Social Averaging: Consensus contains the tendency to shave the specificity off every position until what remains fits through the aperture of collective approval, and what fits through is typically what offends no one and challenges little.

The more perspectives, the worse this becomes, the precise opposite of what diversity-of-voices governance theory promises. Fifty people with different epistemic frameworks and values, averaged, is not a synthesis.

The Münchhausen trilemma

The Münchhausen trilemma is the problem that any justification for a foundational principle either rests on another principle requiring its own justification, circles back to itself, or terminates in an axiomatic assertion that cannot itself be justified. Every governance system, or ethical framework and epistemology meets here.

Power combats power is Madison's solution, and it is the solution built with priors that assume, with fair historical evidence, that you cannot make people good, so you must design systems that function with bad people. Ambition must be made to counteract ambition. It is engineering for the worst case, and it produces something that works the way a hospital works, by managing dysfunction rather than achieving health.

However, there is a way, and it is not a mechanism. It is the patient cultivation of a culture of genuine mutuality. It does not cultivate quickly, and you cannot cultivate it from the outside, nor can you cultivate it using the same logic as the paradigm of which it is demonstrably not.

What makes coercion unnecessary is not a clever mechanism but a unquantifiable quality of social embeddedness, recognition that your conditions and mine are entangled, another word for that, if we are honest, is love, not its ideology or its symbol but its being. When that exists at sufficient depth, members no longer experience their local sovereignty and the collective as opposing forces.

Forests contain disasters; as a whole, the forest is most often in equilibrium. Fire is part of the metabolism of the forest. It creates the conditions for growth that the canopy was suppressing. A forest that never burns accumulates so much dead material that when fire finally comes, it is catastrophic. The forest is not an answer to the trilemma. It is an instance of it. It simply is what

persists, and what persists undercuts and shapes what philosophically is made, by virtue of the stability of its being. It is always in flux, and this ‘admittance’ provides its meta-stability.

The pathology of current governance systems is that instability is suppressed; short-termism contributes heavily. The forest is managed into brittleness by optimising against the process that kept it resilient. Unaddressed grievances, unsampled conditions and manufactured consent sits on the surface, while the actual state of people goes unread.

Equilibrium is a dynamic property of a system that can handle its own disturbances.

DI is a dynamic function of the trilemma. We cannot ground legitimacy absolutely. DI is not a better system. It is an attempt to facilitate the conditions under which the question of what system to have becomes answerable by the people who must live inside the answer.

What Distributed Inference Addresses

As groups grow, voices don’t just get quieter; they get filtered. Social palatability, confidence, narrative fit, and stamina can dominate. The group stops sampling its own intelligence and starts sampling its preexisting status.

By rotating authority, scoping it tightly, and tying it to concrete proposals rather than personalities, DI turns governance into a repeated sampling process. Every cycle is a fresh read of the active signal, rather than a referendum on past social capital.

DI treats governance as a distributed signalling and a consensus acquisition roadblock problem. Each cycle samples the entire active population of ideas, sentiment and energy. Authority rotates to prevent any single social dynamic from becoming self-reinforcing. Even socially awkward, unpopular, or counter-narrative ideas can be tested at low cost through pilots. What persists is what survives repeated, collective contact with reality.

DI posits that most coordination problems are not caused by bad intent but by information distortion. By making proposals, decisions, roles, and outcomes visible by default, DI reduces the surface area for unmerited power to hide. Accountability emerges through selection and non-selection, without judgment. Patterns emerge. Who proposes a lot and never delivers. Who blocks without contributing? Who coordinates well? Reputation is grounded in observable behaviour.

DI suggests that outcomes should persist only as long as they remain subjectively beneficial under current conditions. Stability is achieved through repeated confirmation. What is genuinely good will tend to remain. What requires protection from questioning is already suspect.

DI's claim is not that it produces objectively correct valuations, but that it produces valuations generated by the people actually affected rather than by those with the most to gain from a particular outcome.

Every structural choice in DI is ultimately an answer to the same question: how do we preserve high quality representation in which every member can participate meaningfully, in a psychological distrubiton of people that contains every possible interest that has a motive to corrupt. DINs lateral architecture is a distributed system designed to prevent the single points of failure that centralised governance invariably produces.

The assumption that you can compress the input, that a mandate, a utilitarian calculation, an expert consensus can adequately represent the information content of millions of actual lived conditions without catastrophic loss, is false.

DI aims to create enough agreed formal friction to allow reference and consensus to a proven frame to become the socially grounded foundation that prevents pure power/dependency-based outcomes. By balancing the social environment, the common's health and utility is its own argument.

DI is not a system designed against human nature. It assumes cooperation as the default and builds to protect that default from the specific and identifiable things that corrode it. It is trust-building without being trust-dependent. It creates conditions where trust can emerge through repeated interaction with visible outcomes; it builds epistemological humility into the governance structure itself rather than simply recommending it as an attitude.

Core Principles

Time-bounded, scoped authority to prevent social capture.

Radical transparency to preserve signal integrity and accountability.

Mutability: Mutable rules may change with moderate, participation-scaled pressure. They support adaptation, tuning, and local optimisation.

Semi-immutability: Semi-mutable rules may change only with sustained, high participation and delay. They protect system stability, exit rights, and long-term coherence.

Consent to the frame: People don't need to agree on each decision if they agree on the legitimacy of the decision-making mechanism.

Clear roles: Flat hierarchies sound virtuous but perform poorly. What works is bounded authority: roles with power, scope, accountability, and expiry dates. Authority that is visible and revocable is less dangerous than de facto authority.

Exit must be easy and non-punitive: Nothing poisons consensus faster than hostage dynamics. If leaving is costly or shameful, people fight to control the system instead. Easy exit lowers stakes, reduces domination, and paradoxically increases cooperation.

On **Continuous Provisionality** and **The Semi-Mutuality of Externality**

Every tenet of DI is conditional, adaptable and provisional, it does not know the way because there is no unsubjective way, it lets subjectivity live.

Say someone disagrees with tenets or wants to fork, or merge his ideas with DI. Who am I? The author, to disagree without being an opposing authority? If I use the document's own terms, we'll either merge or fork or just live to disagree, the document is intergrating something other than itself to itself, and avoiding its own closed loop.

If too much change is enacted at once, it destroys the very form that makes DI something worth existing, as something containing its own semi-mutuality of externalities. Thats why it's semi-mutual, the same way a human cell permits most but not all that crosses its membrane.

Progressivism must recognise that it is born of its opposite, and if it quashes its opposite, it becomes its own tyranny, slack to its own change. Progressivism that makes sense to itself produces as much resistance as is corrective, and no more; otherwise, no multidimensional dialectic, no fruit.

A cell in a body is a semipermeable membrane that works for the good of the body because it depends on and is the body, if it is not mutual to the surrounding cells, the good of the whole body ignores, recycles or dismantles that cell, however the cell can also mutate in a way that is beneficial to the whole of the body. In a way, cancer is a hijacking of the body's consent of its own holistic health; cancer is like unchecked cell individualism.

Self-similaritiy at scale

As DI is adaptable and not a set object, it must be multi-modal, to scale to a global level, suiting the populations it serves, fit to where that subjectivity lives, yet self-similar enough to agree on a definition of what a DIM is, to work as a network, and not be classed/be classed as an external system.

It is not that every DIM looks the same, but that every DIM is recognisable to every other DIM as operating by enough shared logic that coordination is possible without a governing layer to certify the resemblance.

So DIN must be simultaneously:

Coherent enough to be a network. Hollow enough to resist being a target. Self-similar enough to recognise its own instances and mutable enough that no instance can claim to be canonical.

Nature demonstrates that optimal global systems does not require global directives, but local responsiveness operating according to shared grammar.

DI is the generative rule. Each DIM is a self-similar instance. DI itself is just the sum of all DIMs.

Eligibility, Participation, and Legitimacy

DI uses “member” pragmatically: anyone recognised by a community as eligible to use the process within that community. Each DIM decides its own eligibility boundary; DI does not impose a universal membership ideology.

Recommended default: eligibility requires sustained presence and/or contribution for a minimum period (X time, or X cycles), and is maintained by a simple member list with social verification.

Exit is always available. Rejoining is permitted, but may include a cooldown (e.g., one cycle) before full voting rights return, to reduce opportunistic churn.

Participation is optional, but legitimacy is not free. DIMs may treat participation rate as a first-class health metric and adapt governance behaviour when turnout drops (see Voting Mechanics).

How a DIM Functions: A Scenario

Note: This example configuration is simple and provisional, the exact configuration(s) are model specific and evolvable.

Community size: 20 active members
Voting frequency: Monthly
Key issue for this cycle: Improving waste management

Step 1: Members Submit Proposals

At the start of the cycle, members submit their ideas for change privately to prevent political lobbying. Proposals can be about anything: implementing projects, changing arrangements, reversing decisions, modifying governance rules, creating or removing mechanisms. The proposal system is unrestricted.

Step 2: Voting Takes Place

Each member gets 3-5 votes (the community defines this) and can distribute them however they like across proposals.

When votes spread across many proposals and none reach majority:

Top N proposals that reach minimum threshold pass. Default: Top 5 proposals reaching minimum 40% threshold. Both conditions must be met: above threshold AND in top N. If fewer than N reach threshold, only those pass. Communities can adjust both N and minimum threshold via proposal.

Example:

10 proposals, 20 voters
Minimum threshold: 40% = 8 votes
Top 5 rule applies

Results:

1. Proposal A: 14 votes → PASSES
2. Proposal B: 12 votes → PASSES
3. Proposal C: 11 votes → PASSES
4. Proposal D: 9 votes → PASSES
5. Proposal E: 8 votes → PASSES
6. Proposal F: 7 votes → FAILS (below threshold) 7-10. All below 7 votes → FAIL

Proposals F-10 can be resubmitted next cycle.

Participation is optional each cycle. Quorum is implicit in thresholds, no enforcement by mandatory voting.

Members may abstain on any proposal without penalty. Abstentions do not count towards thresholds. Failure to vote is treated as abstention.

Any member may flag a formal objection. If objections exceed X percent, proposal auto-converts to pilot, delay, or requires revision.

Participation health: Track participation per cycle (and optionally a rolling average). If participation remains below a defined threshold for Y consecutive cycles, the DIM may enter a stabilisation: fewer irreversible changes, heavier use of pilots, and automatic review of ongoing arrangements.

Step 3: Facilitators Are Selected

The top 5 vote-getting proposers whose proposals passed become Facilitators for this cycle.

Example:

- Build composting system (Hana) - 14 votes → PASSES - Hana is Facilitator
- Organise skill-sharing night (Ben) - 12 votes → PASSES - Ben is Facilitator
- Install solar panels (Dan) - 11 votes → PASSES - Dan is Facilitator

Result: 3 proposals pass. Hana, Ben, Dan are Facilitators for this cycle.

What Facilitators Do

Facilitators coordinate implementation of passed proposals.

Their role:

- Oversee implementation of their proposal
- Coordinate with other facilitators
- Run cycle meetings
- Track and report progress

Duration: One cycle only. Eligibility resets after cycle ends. Can be selected again if future proposals pass.

Selection is natural: Repeated facilitation demonstrates competence. Proven facilitators become Representatives for cross-DIM coordination. Competence is rewarded through continued selection. Poor performance leads to proposals not passing.

Step 4: Execution & Accountability

Hana coordinates the composting system implementation. Ben organises the skill-sharing night. Dan coordinates solar panel installation. At the end of the cycle, they document progress for transparency.

Step 5: Review & Next Cycle

If a proposal doesn't work, members can vote to modify or remove it next cycle. If it succeeds, it becomes part of the community's evolving system.

Proposal Scope and Formalisation

Not everything needs formalisation. Daily life requires fluid agreements between individuals and small groups.

Proposals are unrestricted. You can propose implementing projects, changing arrangements, reversing decisions, changing rules, creating mechanisms, removing mechanisms—anything. The proposal system contains all possible governance adaptations because it's radically open to its own transformation.

The boundary between informal and formal is itself governed by the community.

Default principle: Act unless someone objects, then formalise.

Any member may propose that a category of decision requires formal proposals. Any member may propose that a category no longer requires formalisation. These scope-change proposals follow normal voting rules and thresholds.

Bootstrap minimum for new communities:

- Changes to shared physical spaces or resources
- Allocation of community funds or collective property
- Changes to explicit community norms or agreements
- Changes to the governance system itself
- Any decision where at least one member requests formalisation

This initial scope is temporary scaffolding. After the first cycles, the community may propose changes to scope itself.

Adaptive scope: Over time, high-trust communities with easy coordination may create fewer formal proposals and prerequisites. Communities experiencing conflict or coordination issues may expand formalisation to fit their challenges.

In areas where communities have not yet formalised, existing local bodies can hold proxy membership on behalf of communities that haven't yet entered the network.

The proxy is redundant when the community develops enough to represent itself. It creates no obligation on the community being proxied.

Existing local institutions that adopt DIM as a pilot aren't running the DIM; they enable it. If enough people engage, the DIM forms, and the institution's proxy role ends.

A DIM can take on any name and aesthetic; it belongs entirely to the community using it.

Representation Overhead

Any act of governance is an act of measurement, and measurement has costs. Increasing the fidelity of representation increases those costs faster than the fidelity gained, and more data requires more governance to interpret, which requires more data to govern well. A situation where a community is primarily engaged in representing itself is self-defeating.

DI aims to condense social capital below a point where this is problematic. The proposal system can dismantle mechanisms as readily as it creates them.

Scarce Shared Resources

If a proposal allocates scarce shared resources above a DIM-defined threshold (money, land, housing, critical infrastructure time), it should explicitly include: (a) what is being allocated, (b) who loses access or opportunity, (c) sunset/review conditions, and (d) fallback if the proposal fails or underperforms.

Ongoing Arrangements vs Time-Bound Work

Some proposals create time-bound implementation work: "Build a composting system," "Organise a skill-sharing event," "Repair the solar panels." The facilitator role exists only for the duration needed to complete the work.

Other proposals establish ongoing responsibilities: "Alice manages waste collection," "Ben coordinates tool maintenance, "Establish a weekly community meal rotation." These arrangements continue indefinitely once approved. They do not need to be re-proposed each cycle. The default state is continuation unless someone submits a proposal to change the arrangement.

Any ongoing arrangement can be challenged at any cycle: "Rotate waste management to Cara," "Split tool maintenance between Ben and Dan," "End the weekly meal rotation." If the challenge wins sufficient votes, the arrangement changes. If not, it continues unchanged.

Stable, working arrangements persist without administrative overhead. Poor performance or changing needs can be addressed quickly without waiting for fixed-term limits.

Obligations, Consent, and Non-Compliance

Obligations within a DIM only arise through **explicit consent**. An obligation exists if and only if a participant has:

- Voluntarily accepted a role,
- Agreed to an ongoing arrangement,
- Opted into a task, responsibility, or resource usage with stated conditions.

No member is obligated by default to labour, participation, or compliance beyond the degree they have explicitly accepted.

The default response to non-compliance follows a simple escalation:

1. **Notice** – the unmet obligation is identified and communicated.

2. **Scope-limited restriction** – loss of access or privileges directly tied to the obligation (eg. stepping down from a role, loss of access to a resource governed by that role).
3. **Process-out** – removal of participation rights within the DIM if non-compliance persists or undermines shared functioning.

Sanctions are limited to **procedural consequences** (loss of access, role removal, or participation rights) and do not extend to punishment, asset seizure, or personal retaliation. Participants may always exit an obligation by resigning a role or withdrawing from an arrangement, subject only to any explicitly agreed transition conditions.

Re-confirmation and Anti-Entrenchment

Because ongoing arrangements may continue by default, DIMs may add a lightweight re-confirmation cadence to prevent silent tenure. Example: every N cycles, ongoing roles are re-confirmed, auto-renew unless objections exceed X%, or trigger a confidence check vote.

For high-leverage ongoing roles (finance, legal interface, safety), redundancy is recommended: deputies/shadowing, handover notes, and mandatory rotation or cooling-off periods.

Easy Pilot Framework

To enable innovation and experimental proposals that may not have majority support:

Any proposal can be framed as a pilot with limited timeframe, limited scope, defined evaluation criteria, and default sunset (expires unless renewed).

Lower threshold for pilots: Default 30-40% instead of 50%, because risk is bounded, changes are reversible, and pilots generate learning.

Example: Instead of "Implement permanent composting system community-wide," frame as "Try composting system as a 2-month pilot in Zone A, evaluate after."

After the pilot, results are presented to the community. Proposal to continue/expand requires a standard threshold. Can sunset without stigma.

This preserves requisite variety, the ability to try unpopular ideas that might turn out to be good.

Proposal Limits

To prevent proposal spam while maintaining openness:

Default: 2-3 proposals per person per cycle. Communities can adjust via proposal.

Natural disincentive: Proposer of passed proposal becomes Facilitator. The facilitator role is work. Submitting frivolous proposals risks unwanted facilitation responsibility.

Social cost: Repeated low-quality proposals damage reputation. Future proposals less likely to pass.

Representatives and Cross-DIM Coordination

Representatives coordinate between DIMs when proposals affect multiple communities or require cross-DIM cooperation.

Selection: Proven facilitators from recent cycles (default: 3-5 most recent). Facilitators may collectively decide who represents them. If no agreement, the most experienced facilitator. Maximum term: 2-3 cycles before mandatory rotation.

Role: Carry information between DIMs. Negotiate coordination on behalf of their DIM (within bounds set by their DIM). Bring cross-DIM proposals back to their DIM for vote. Cannot make binding commitments without DIM approval.

Authority boundaries:

- Can commit to: information sharing, coordination timing, procedural details
- Must return to DIM for: resource allocation, binding agreements, major decisions

Representatives coordinate but don't govern.

Passing a proposal authorises coordination and commitment within the scope of that DIM. Implementation still depends on real capacity, and cross-DIM proposals require approval from each affected DIM. A proposal can pass and still fail socially or structurally if the required cooperation does not materialise.

Two Types of DIMs

DIN consists of two types of DIMs with different remits:

Geographic DIMs are based on location/proximity. They govern shared physical space, local resources, neighbourhood norms. Examples: neighbourhood, housing cooperative, village, district.

Functional DIMs are based on shared service or function. They cross-cut geographic boundaries. People participate based on relevance. Examples: healthcare service, education network, infrastructure maintenance, care coordination.

Functional asset DIMs present a genuine exception. A specialist operator running a server farm, energy installation, or cargo vessel cannot rotate out on the standard facilitation cycle without operational cost to everyone, depending on the asset. Specialist tenure is therefore operationally continuous, subject to the same confidence review and deputy requirements as other roles.

Individuals can be in multiple DIMs (one geographic plus several functional). You vote in each DIM only on proposals relevant to that domain. Engagement varies by DIM—active in geographic, occasional in functional.

When geographic and functional DIMs conflict: The geographic DIM can choose not to engage that service. Functional DIM can choose not to serve that geography. Or: negotiate interface rules through Representatives.

Cross-DIM Coordination Mechanisms

DIMs are sovereign but often need to coordinate. This happens through Representatives.

Simple coordination: Representatives meet and communicate directly. Work out details. Report back to their DIMs. Binding commitments require DIM vote.

Proposals affecting multiple DIMs: Originating DIM's Representative brings proposal to affected DIMs. Each DIM votes independently. All affected DIMs must approve for proposal to proceed. If any disagree: negotiation, compensation, or stagnation.

Conflict resolution when DIMs disagree on resource allocation or proposals with ripple effects:

Compensation: Majority DIMs can proceed but must compensate affected minority DIMs. Terms negotiated by Representatives. Both sides must approve compensation.

Exit: Minority DIM(s) exit the cooperation. Find alternatives. No forced participation.

Stagnation is a legitimate outcome. If DIMs cannot reach agreement, the proposal fails. This is preferable to forcing compliance. Indicates cooperation model doesn't work for this issue. Drives better compromise or separation.

By the time complex coordination is needed, such as 6 DIMs sharing a watershed facing major resource allocation decisions—they'll have history with each other, established norms for negotiation, proven Representatives, trust or lack thereof. Context-specific solutions emerge from their actual relationship.

Late-stage DIN coordination:

Inter-DIM proposals are not a single category. Three tiers have different legitimacy requirements.

Tier One is bilateral: two DIMs agree to coordinate on something that primarily affects only them. Any DIM may enter bilateral agreements with any other DIM. The only requirements are that the agreement passes at elevated internal threshold (60% recommended, adjustable by proposal) to reflect that external obligations warrant stronger internal consensus than internal arrangements, and that the agreement is visible on the network by default.

Tier Two is regional: a proposal affecting a defined cluster of DIMs sharing geography, infrastructure, or a resource system. The originating DIM's Representative proposes a list of affected DIMs. Any DIM may claim affected status within a notification window (default 30 days) before voting proceeds. Failure to notify an affected DIM does not invalidate the agreement but does entitle that DIM to a review cycle. Regional agreements require approval from all affected parties. DIMs at elevated thresholds. A DIM that disagrees may exit the proposed arrangement; stagnation is a legitimate outcome.

Tier Three is Network-wide: proposals that establish norms, protocols, or constraints intended to apply across an entire DIN or large federation. These require a constitutional threshold: supermajority approval (75% recommended) from participating DIMs whose combined membership exceeds a defined floor (suggested: 40% of total active DIN membership). A collection

of small or newly-formed DIMs cannot establish network-wide norms. Scale of adoption confers legitimacy.

What makes a DIM affected: a DIM is affected if the proposal, if implemented, would materially change its access to shared resources, its obligations to other DIMs, or its members' safety conditions, regardless of whether it was named in the originating proposal.

A DIM of 20 people can make bilateral treaties freely. It cannot unilaterally bind the wider network.

Boundary Constraints and Non-Negotiables

Some issues create irreversible externalities across DIM boundaries (e.g., pollution, safety hazards, shared watershed impacts). DIMs may define “boundary constraints”: non-negotiable limits on actions that impose irreversible harm on others.

Boundary constraints are adopted through explicit cross-DIM agreement (treaty-style proposals) and are enforced primarily through refusal of cooperation, withdrawal, and exit rather than forced compliance.

Scaling Through Mesh

Traditional scaling works through nested layers: Local → Regional → National. Each layer governs the layer below, creating hierarchy and concentrating power at the top.

DIN scaling is a lateral mesh. No super-governing layer. Representatives communicate laterally. Coordination is negotiated. No DIM governs other DIMs. Disagreement means stagnation or separation.

At 200+ people within single DIM: Chunk into 3-5 zones. Vote on proposals closer to you. Can vote on other zones if you choose. Top 5 facilitators coordinate across zones.

Beyond 200, don't nest, fork: Create multiple DIMs. Connected through Representatives. Lateral communication. No governing DIM above them.

At large scale: Web of DIMs, not a pyramid. Geographic DIMs (place-based) and Functional DIMs (service-based). Representatives negotiate at boundaries. Distributed Inference Network.

How does DIN build a railway between them if they have to negotiate with hundreds of sovereign DIMs along the projected route? How can a functional legitimate collective override of local veto without reintroducing the centralised authority it is designed to replace?

If you are inside DIN, your sovereignty is an artefact of collective agreement. The collective can, through prior agreement and before any specific railway is proposed, specify the conditions under which local veto over collective necessity is and is not a legitimate exercise of that sovereignty. Pre-commitment before specificity. The Tier Three mechanism is also pertinent here, but much remains unknown, and issues remain. It is an inherent tradeoff of a decentralised system.

A DIM in one locality doesn’t wait for a message from another. They are subscribed to the Network. Data is held in a shared commons. A DIM doesn't fundraise for the design of a new solar array.

They download the design from the global Network. The funding is a local matter of allocating parts and labour from neighbouring DIMs.

Knowledge that is open becomes the commons, live and updating.

Voting Mechanics

The requirements for any voting system: votes must be verifiable (you are who you say you are), anonymous (can't see how individuals voted), auditable (can verify count was correct), and accessible (people can actually vote).

Different DIMs will need different implementations based on size, tech access, threat model, and cultural norms.

We don't prescribe the solution. Your voting mechanism must achieve these properties.

One approach that balances trade-offs:

In-person meeting component for freeform discussion, group decisions for non-contentious matters, proposal refinement, and non-binding "temperature checks."

Actual voting through paper ballots during meeting (for those present) and digital voting window (48 hours, for those who couldn't attend). Identity verification through simple passphrase plus social verification (you're on the member list, your community vouches for you). Anonymous ballots (identity verified but vote is private). Public results (total counts visible to all).

This balances accessibility, verification, anonymity, and simplicity.

High-Stakes Proposal Classes

High-stakes proposals include, but are not limited to:

- Allocation of significant shared funds or property,
- Legal or financial commitments binding the DIM,
- Changes to governance structure or voting rules,
- Long-term exclusion or safety restrictions,
- Actions with irreversible cross-DIM externalities.

High-stakes proposals may require higher participation thresholds, longer delays, pilot defaults, or additional safeguards.

Turnout Floors and Stabilisation

Thresholds (e.g., 40% + Top N) protect against fragmentation among active voters, but do not guarantee broad legitimacy under low participation. DIMs may therefore optionally define turnout floors, particularly for high-stakes categories (funds, property, legal commitments, governance changes).

If turnout falls below X%: options include (a) only pilots can pass, (b) only reversible changes can pass, or (c) all proposals auto-delay. The chosen rule is itself subject to proposal.

Participation health can be treated as a rolling metric. If turnout remains below X for Y consecutive cycles, the DIM may enter a stabilisation mode: increased use of pilots, higher thresholds for irreversible commitments, and scheduled review of ongoing arrangements.

On Weighting Mechanisms

Aggregation of votes is not perfectly uniform. Small, symmetric variation is introduced at the point of accumulation to prevent exact vote count manipulation and pure social capital from overreaching. Expected outcomes remain as such, but exact rankings are not rigid.

An earlier version of this architecture included a full vote-weighting protocol: repeated cycle wins applied a diminishing weight to future proposals, persistent low-vote performers received a compensatory buff, and a separate voting sheet distinguished mutable from immutable categories with their own thresholds. The intent was to prevent social capital dominance so that less popular yet robust ideas to come to light.

It was abandoned. The mechanism was its own problem. A debuff for effectiveness is a penalty for competence. A buff for persistent low performance potentiates reverse incentives and sandbaging. A separate immutable voting sheet requires a prior decision about what counts as immutable, which is itself a governance question requiring a vote, and so on. The protocol generated more administrative surface than the problem it was solving, and in doing so, demonstrated the argument it was trying to prevent.

Participation Subsidies

To counter time-wealth bias, DIMs may allocate real money for participation subsidies: childcare, transport, accessibility needs, translation, facilitation time, and other costs that otherwise concentrate governance in the time-rich. These subsidies are scoped, auditable, and tied to reducing participation friction rather than policing motivation.

Transparency and Privacy

Transparency by default applies to the governance surface: proposals, outcomes, role assignments, and execution reports. Personal disputes, sensitive personal data, and private communications remain private by default.

Record-keeping may be constrained to: proposal text, vote totals, role assignments, and DIM structure changes. Higher-stakes decisions may warrant stricter verification and stronger anti-coercion norms (e.g., no vote receipts, no mechanisms that allow individuals to prove how they voted).

Legal and Financial Interface

DIM/DIN operates on principles incompatible with most external legal systems. A translation layer is required.

Trust/Charity structure: Assets held by a legal entity (nonprofit, cooperative, LLC). DIM governs the entity through its proposals. Officers rotate but the entity persists.

Multi-signature financial controls: Bank accounts require 3+ facilitators to authorise transactions. No single point of financial control. Transparent to DIM members.

Cryptocurrency options: Multi-sig wallets requiring multiple facilitators. On-chain transparency. Cannot be frozen by external authority.

As DIN becomes economically significant, DIN-to-DIN transactions bypass external systems. Internal currencies and mutual credit emerge. Legal systems may adapt to recognise DIN entities. Traditional methods gradually replaced by DIN-native methods.

Web3 infrastructure might solve specific DIN problems. Multi-signature wallets requiring multiple facilitators to authorise transactions remove single points of financial control without requiring a bank. On-chain transparency makes financial balance visible by default. Smart contracts can automate resource mobility. If DIN becomes economically significant, DIN-to-DIN transactions can bypass legacy financial systems entirely through internal currencies and mutual credit built on existing Web3.

Funding and Anti-Capture Defaults

A single wealthy participant can attempt capture through dependency: funding essential needs, controlling assets, or attaching covenants to external contracts. DIMs may adopt simple funding defaults to reduce this risk.

Recommended defaults include: caps on any single donor/creditor as a share of the annual budget; acceptance of large contributions only in standardised forms (unrestricted donation, narrowly restricted donation, or loan with no governance covenants), and prohibitions on personal guarantees and private ownership of choke-point assets used by the community.

Membrane roles (bank access, signatories, legal interface) should be separated from major donor/creditor status, rotated, and bounded by explicit authority limits. External commitments that bind the DIM or its assets should require explicit DIM approval.

Interface with Traditional Systems

Some functions require hierarchical coordination: military (speed, secrecy, chain of command), emergency services (immediate response), critical infrastructure (cannot fail), international diplomacy (stable negotiating partners).

When hierarchy makes decisions affecting DIN, information comes through Representatives. DIN votes on response/compliance. Decision is local—this DIM's relationship to that hierarchy.

Example: Infrastructure - Regional utility (hierarchical). Local DIM votes on: acceptance of service, payment terms, local standards. Can exit and create alternatives if possible. Or negotiate through Representatives.

DIN handles what can be handled locally/voluntarily. Hierarchy handles what requires speed/enforcement/scale. Interface is explicit negotiation. As DIN grows, more functions can be handled DIN-natively. But some hierarchical functions may persist indefinitely.

Economic Coordination and Value Exchange

DI does not prescribe a single economic ideology. It operates as a base upon which multiple value-exchange mechanisms may emerge and coexist.

At small scales, particularly within geographic DIMs where members share daily life and high-trust relationships, coordination may occur without formal currency. Labour is often motivated by shared identity, mutual care, reputation, and direct reciprocity. In such contexts, explicit accounting is frequently unnecessary and may even be counterproductive.

As coordination extends across DIMs, distance increases. Emotional salience weakens, time horizons lengthen, and participants may not share a lived context. In these conditions, informal reciprocity alone becomes insufficient. Cross-DIM coordination requires portable signals of contribution and obligation that can persist beyond immediate social memory.

These signals may take many forms: time-based credits, mutual credit systems, scoped internal currencies, reputation-weighted contribution records, or other mechanisms developed locally. Regardless of form, effective value-exchange mechanisms share common constraints:

- Contributions must be readable beyond the immediate context in which they occurred.
- Accumulation must be bounded in scope, duration, or convertibility to prevent extractive dynamics.
- Credits represent contribution, instead of ownership or control.
- Exchange mechanisms are domain-specific where possible, rather than universal and fungible.
- Value systems remain subject to proposal, modification, and removal through normal DIN processes.

DI distinguishes between exchange and accumulation. Fair value exchange is compatible with DIN at all scales. Unbounded accumulation divorced from contribution is not assumed, required, or protected by the system.

Economic mechanisms within DIN are expected to evolve. Early-stage communities may rely primarily on intrinsic motivation and informal reciprocity. As networks grow and interdependence increases, more explicit accounting may emerge to support complex coordination. No mechanism is permanent; all remain provisional and revisable.

DI does not attempt to eliminate self-interest incentives. Instead, it aims to increase accuracy and widen self-interest incentives by tightening feedback loops between contribution, reward, and collective impact. As adoption grows, value exchange, in theory, becomes increasingly contextual, relational, and visible, reducing the incentive for extortion.

DI is not ideologically opposed to concentrated wealth. Extreme imbalance is partly a symptom of broken representation. When people are genuinely represented, and their actual conditions are sampled rather than their manufactured consent, arrangements that serve the few at the expense of the many don't get vouched for. They are typically not what people choose when they have an informed say.

Private Property, Assets, Use-Rights

DI does not require abolition of private property. It distinguishes fair value exchange from extraction dynamics.

Recommended framing: (1) personal property is protected; (2) private productive property is permitted where it does not create dependency leverage over others; (3) commons property is governed by DIM proposals; (4) hybrid arrangements may exist where private title is paired with community use-rights.

Extraction risk increases when a private arrangement creates dependence (housing, essential infrastructure), enables rent without ongoing contribution, or concentrates choke-point control. DIMs may adopt review triggers, right-of-first-refusal, and conversion/buyout paths for assets that become essential to community functioning.

Large-scale productive assets, energy installations, server infrastructure, cargo vessels, manufacturing facilities, require operational hierarchy. A working oil rig needs chain of command and the capacity for immediate unilateral decision under conditions where delay is dangerous or fatal. DI leaves internal schema mostly alone. These are questions with operational answers and the answers differ legitimately by context, risk profile, and the nature of the work itself.

DI concerns itself with the asset's relationship to the network. How the communities depending on it maintain meaningful accountability and how the people working within it remain within the wider network of mutuality.

DI is compatible with hierarchy where hierarchy serves the work. A rig can have a captain. The captain's authority ends at the DIN's edge. The Geo-DIMs receiving the rig's output govern that relationship through their own processes.

Private ownership of productive assets is compatible with DIN. Personal property is protected. Private productive assets at any scale are permitted. DIN prescribes nothing about ownership structure, internal organisation, or what return a private owner considers worthwhile.

DIN is a voluntary cooperative; DI replaces compulsion-based incentives with consequence-based ones. The current system says participate or be sanctioned. DI suggests, participate and receive value, don't participate, and it moves around you.

A privately owned asset that wishes to exchange with the DIN does so on agreed terms or not at all. This requires no regulation, no oversight body, no enforcement mechanism. DIN collectively declines exchanges it hasn't agreed to, in exactly the same way any individual DIM declines coordination with actors whose terms it finds unacceptable. This is the exit right.

Agreed terms mean a treaty between the asset's owners and the relevant DIN participants, negotiated openly, visible by default, covering margin, surplus distribution, maintenance obligations, and review conditions. The treaty is the precondition for participation. Without it, the asset sells elsewhere or operates externally.

Where a private asset has already created dependency before treaty terms are established, DIN's collective exit right becomes practically difficult to exercise. The rut of negotiating from inside a dependency you are trying to exit. DIN’s interest is in establishing treaty terms before.

What constitutes legitimate margin and what constitutes unjust profiteering is defined through treaty negotiation between the asset owner and the DIM. If the DIM finds the proposed margin excessive, it declines. If the owner finds the DIMs counterterms unacceptable, they sell elsewhere. The treaty is the market between a private owner who wants access and a DIM that can collectively choose whether to provide it.

The DIM's collective judgment about acceptable terms is the only regulator required. It reflects the actual conditions of the actual people affected, rather than an external definition of fairness.

Distributed Technology and Sovereignty

Solar, battery storage, rainwater filtration, community growing, open source manufacturing, distributed fabrication. These are plausible at community scale, at accessible costs and are getting cheaper faster than centralised alternatives. A DIM only requires enough internal resources that the floor is met by DIN before it interfaces with external systems. It maintains the exit right at the material level.

Open source knowledge commons helps to facilitate this point. A DIM could pull from a global hub of documented solutions, tested implementations, locally adapted variants.

Distributed infrastructure and distributed governance are becoming the same argument.

Safeguarding and Grievance Track

DIN does not rely on ad hominem proposals for interpersonal conflict. Safety and grievance handling runs on a separate track designed to be fast, bounded, and minimally capture-prone.

Two-Track Structure

Safety track (fast): temporary, bounded measures that prevent immediate harm and reduce escalation.

Resolution track (slower): mediation and, if necessary, a community-backed process for longer-term constraints.

Safety Steward (Rotating) and Second-Key Backstop

A rotating Safety Steward receives confidential reports and can enact time-limited safety measures: separation at shared events/spaces, temporary no-contact rules in shared spaces, and temporary suspension of a specific privilege. Measures auto-expire (e.g., 7–14 days) unless renewed through defined process.

To reduce unilateral abuse, restrictive measures beyond a short window (e.g., >48 hours) require a second key: a co-signer (Safety Witness) who confirms scope and duration. Safety roles rotate and are subject to re-confirmation like other ongoing arrangements.

Mediation First

Where safe, voluntary mediation is attempted before escalation. Mediation may be internal (a rotating mediator) or external (preferred where local bias is high). Agreements can remain informal or be formalised as an ongoing arrangement.

Sortition Review Panel (Only If Needed)

If mediation fails or sustained restrictions are requested, a temporary review panel can be formed by random selection (sortition) from eligible members, excluding involved parties and clear conflicts of interest. The panel reviews evidence privately and recommends a time-bounded outcome.

Long-term restrictions (multi-cycle exclusion, revocation of participation access) are governance constraints and therefore require formalisation through the proposal system under the DIM's high-stakes rules.

Privacy and Transparency

Grievance content is private by default. A minimal public meta-log may record that a safety action occurred, its type, and its expiry, without exposing sensitive personal details.

Known Practical Limitations

Acute crisis requiring immediate unified action: Rolling cycles are too slow for true emergencies. Facilitators can make urgent decisions, but if a crisis requires sustained centralised coordination beyond a few days, DIN may need to temporarily suspend normal operations or defer to specialised emergency systems.

Highly technical domains requiring deep specialisation: When 95% of members lack knowledge to meaningfully evaluate proposals in a domain, domain trust can concentrate into de facto permanent authority.

Groups below ~15 members: Too small for meaningful rotation. Everyone knows everyone's capabilities intimately. Formal mechanisms add overhead without benefit. DIN is optimised for 20+ members.

Groups above ~200 members without subsidiarity: Direct participation becomes unwieldy. Zone chunking required, introducing new coordination costs and potential capture points.

When exit isn't possible: If members can't actually leave (economic dependence, geographic isolation, family ties), pressure builds. DIN assumes exit.

Ideological capture during formation: If the initial seed group is ideologically homogeneous and fork/variation mechanisms aren't established early, DIN can become just another rigid or dogmatic system.

What Distributed Inference Cannot Do:

Resolve fundamental value conflicts. If members have irreconcilable core values, no process makes them compatible.

Create trust where none exists. DIN can protect existing trust and slow its degradation. It cannot manufacture it.

Eliminate human status games. It can only make them more open and less destructive.

Function without minimum participation. If 70% of members fully disengage, the system reduces to minority rule.

Force genuine cooperation. Formal approval doesn't guarantee social support.

Social Reality and Formal Mechanisms

Formal governance can create accountability structures, decision procedures, and transparency. It cannot force genuine cooperation, prevent social dynamics, or resolve fundamental disagreements.

A proposal can pass formally but fail socially. If not enough people actually help, resources are "unavailable," or there's passive resistance, implementation fails regardless of formal approval.

What DI provides: Visibility (failure isn’t hidden), accountability (clear who voted, who helped, who didn't), learning (can propose meta-solutions to coordination problems).

Oscillation: Proposals and counter-proposals can create seesaw dynamics. This arrives at homeostasis through the cost of instability becoming visible. People propose longer windows, higher thresholds for reversals, and pilot periods. Social cost against constant reversal. This is self-regulated through experience.

Informal power persists. Charisma, articulateness, social connection create influence that formal systems can't eliminate. DI makes these patterns visible, but they can and will exist.

Poor ideas can win if enough people vote for them. Democracy doesn't make people wise.

Apathy enables minority rule. If only 30% vote regularly, 16% of the total membership controls decisions. No formal system fixes widespread disengagement.

Mitigations are possible without coercion: treat participation rate as a health metric, use turnout floors for high-stakes decisions, and subsidise participation costs (childcare, transport, accessibility) so governance does not default to the time-rich.

These limitations are inherent to human coordination. The alternative trades these problems for concentration of power, abuse potential, ossification, and reduced adaptation.

Evolution and the Base Model

DIN provides a base (core architecture that must exist), seed configuration (starting parameters that enable viability), and evolution mechanisms (ways to adapt based on reality).

The base:

- Vote on proposals

- Threshold for passage
- Facilitators from passed proposals
- Transparency by default
- Exit always available

The seed configuration (all evolvable):

- Monthly cycles
- 40% threshold
- Top 5 facilitators
- 2-3 proposals per person per cycle
- In-person meeting required

Different DIMs discover different optimal configurations. Successful adaptations spread through observation and forking. Failed adaptations dissolve or iterate. No single "correct" configuration exists. Natural selection operates at the governance level.

Upsides of Current Systems

Conventional hierarchy exists because it solves real problems: speed (someone decides, everyone follows), simplicity (clear chain of command), scalability (nested layers handle complexity), enforcement (centralised power makes defection costly), stability (changes are slow and controlled).

DIN trades these benefits for different ones: participation vs delegation, adaptability vs rigidity, visibility vs obfuscation, local autonomy vs central control, exit vs enforcement.

This makes DIN appropriate for communities valuing these tradeoffs, contexts where costs are bearable, people who prefer DIN's costs to traditional governance costs.

Inappropriate for contexts requiring speed above all, situations needing enforceable coordination, settings where stability is paramount or for people who prefer delegation to participation.

Different tools for different contexts. Given widespread unhappiness with current systems, economic precarity and political alienation, rise of radicalisation, institutional trust at historic lows, DI or some analogue of it is likely to find appeal.

Path to Adoption

DIN cannot be imposed from above. It must prove itself through grassroots emergence.

Phase 1 (now - 10 years): 20-50 person communities try it. Document outcomes honestly. Build a track record. Develop practitioner culture.

Phase 2 (10-30 years): Neighbouring communities see it working (or not). Adopt because it demonstrates value. Network effects emerge.

Phase 3 (30-50 years): DIN communities become viable alternatives. Demonstrate superior outcomes for participants. Younger generations raised in DIN contexts. Political pressure for legal accommodation.

Phase 4 (50+ years): DIN handles an increasing share of actual coordination. Official government becomes increasingly nominal.

DI (or some analogue) may not emerge from institutions. It may emerge through grassroots adoption that becomes impossible to ignore, through demonstrated outcomes that institutional frameworks cannot replicate.

Open Source

DI is open source by design. Its complete specifications are public. Anyone can implement, modify, or fork. No licensing restrictions.

DIM-Compatible Manifest

To reduce confusion and reputational contagion across forks, implementations may publish a simple "DIM-compatible" manifest: which base properties are implemented (proposal voting + thresholds, facilitator selection, role rotation/revocation, transparency surface, exit rights, cross-DIM coordination constraints), and which parameters differ from the seed DIM.

A DIM-compatible implementation need not call itself DIM or use any of the terminology in this document.

On Incompleteness

Distributed Inference is not a complete solution to human coordination. No such system exists or can exist.

The problems DI addresses: power concentration, authority capture, coordination failure, and trust erosion are not fully solvable. They are inherent to human social dynamics.

DI represents a trade-off between simplicity, flexibility, robustness, and honesty.

Complete specification, universal applicability, guaranteed fairness, and immunity to bad faith have been sacrificed. These sacrifices are acknowledgements of theoretical limits. Attempting to eliminate them would generate the bureaucratic bloat and authoritarian creep that DIN exists to prevent.

The system works when members want it to work. This is the inescapable foundation of all cooperative systems. Mechanics can make cooperation easier, can make defection costlier, and can make power visible. They cannot manufacture willingness to cooperate.

When cooperation breaks down despite structural support, the system's response is visibility (making the breakdown explicit), adaptation (changing what isn't working), and exit (allowing incompatible members or groups to separate).

DI is meta-stable, not stable. It maintains itself through continuous adjustment. The goal is to remain capable of seeking a new equilibrium as conditions change.The choice is between honest incompleteness and false completeness that masks dysfunction.