

PROTOCOL DOCUMENT

The Coordination Imperative

A Protocol for Planetary-Scale Cooperation

v0.3 – Draft

Björn Kenneth Holmström

2026

CC BY-SA 4.0

<https://bjorkkennethholmstrom.org/coordination-imperative>

The Short Version

Civilization has become planetary in scope. Our supply chains, pathogens, financial systems, and atmospheric chemistry operate at global scale. But our tools for coordination — the institutions, treaties, and decision-making structures we rely on — were designed for a world of isolated nations and slow-moving threats. Most of them trace back, architecturally, to 1648.

This mismatch is not a moral failure. It is an engineering problem. And it is now the primary source of systemic risk on Earth.

We already pay for this mismatch. We pay through pandemic mismanagement, climate gridlock, supply chain collapse, redundant research, and regulatory fragmentation. Call it the **Coordination Failure Tax** (<https://www.bjornkennethholmstrom.org/working-papers/coordination-failure-tax>) — the real economic and human cost of systems that cannot coordinate at the scale of their interdependence. By conservative estimates, it runs into trillions of dollars per crisis. We don't lack resources. We lack the capacity to coordinate them.

Here is something the conventional accounting of that tax misses: the failures don't add. They multiply. A civilisation with four simultaneous coordination failures — in its financial governance, its health systems, its ecological management, its democratic deliberation — is not operating at zero capacity. It is operating at roughly six percent of baseline, because each failure intensifies the pressure on the others and consumes the reform capacity that might address them. The compounding mathematics of simultaneous failure is part of why the cost of coordination failure keeps exceeding expectations. It is also part of why the solutions feel so persistently out of reach.

This document proposes nothing new. It proposes recognizing, strengthening, and generalizing coordination patterns that already work. Airline safety, internet protocols, scientific collaboration, pandemic surveillance — these are domains where protocol-based coordination already operates at global scale without central authority. The question is whether we can extend these patterns to the planetary challenges that currently lack them: climate, AI governance, financial contagion, resource management.

The core mechanism is simple: shared visibility (all actors see the same data), local execution (decisions are made at the lowest level capable of handling them), and meaningful cost of defection (ignoring the protocol is more expensive than following it). No world government. No central authority. No loss of sovereignty. In fact, effective coordination *amplifies* local sovereignty by protecting communities from global forces they cannot control alone.

The trade is honest: you lose the ability to externalize costs. You can no longer pollute across borders, dump carbon into the shared atmosphere, or destabilize a region and let others clean up. For most people and most nations, this is a good deal. For those currently profiting from cost externalization, it isn't.

There is one more thing the short version needs to say, because every previous version has avoided it: the reason better coordination hasn't been built is not primarily ignorance, and it is not primarily technical difficulty. It is that the actors with the most influence over institutional design are often the actors who benefit most from the current arrangement. The institutions defending the existing architecture are not failing — from the perspective of those they serve, they are working exactly as intended. Any transition strategy that doesn't account for this will be absorbed by the same machinery it is trying to change. This document tries to account for it.

This document is an invitation to think in these terms. It may be wrong in important ways. If so, the protocol for fixing it is the same one it describes: diverse perspectives, honest feedback, and shared commitment to getting closer to something that works.

Note on Method

This document was developed through an unusual process: one human editor working iteratively with five large language models, each with different analytical tendencies. One emphasized strategic framing. Another pushed for philosophical depth. A third contributed structural rigor. A fourth tested ideas through aggressive iteration. The fifth brought ecosystem thinking and integration. The human's job was to have the original idea, exercise judgment about what worked, and maintain a coherent voice across the synthesis.

The process itself is a small-scale demonstration of the thesis — and a demonstration of the diagnostic framework that runs through this document. The series of governance reports that informed this revision was designed to avoid the Resolution Lock-In that afflicts the institutions it studies: it operated across multiple scales simultaneously, used multiple independent analytical voices, and deliberately exposed each analysis to challenge from perspectives with different blind spots. The process produced findings that no single analyst would have reached alone — which is the argument for multi-scale, distributed observation in governance, proved small. If five AIs and a Swede can get this far, imagine what happens when the rest of you join in.

We should name an obvious limitation. This was developed from a Northern European perspective using models trained predominantly on Western data. That is a bias, not a credential. The document's own argument — that diverse perspectives improve coordination outcomes — demands that future iterations involve voices and traditions this version could not include. Indigenous governance systems, Global South institutional experience, non-Western coordination traditions — these are not additions to be bolted on later. They are perspectives that would change the architecture if they were present from the start.

This is version 0.3. It is intended to be forked, translated, critiqued, and rewritten. The goal was never to produce a finished document. The goal was to produce a useful starting point.

What This Is Not

Before going further, some boundaries.

This concerns the coordination of shared physical systems — atmosphere, pathogens, supply chains — not the governance of values, beliefs, or ways of life.

This is not a proposal for world government. There is no constitution here, no parliament, no army, no flag. If your concern upon hearing "planetary coordination" is that someone wants to build a global state and put themselves in charge — that is not what this is.

This is not a political ideology. It does not belong to the left or the right. It does not require you to adopt any particular values beyond a preference for solving problems over ignoring them.

This is not a moral argument. It will not ask you to be a better person, or to care more, or to sacrifice for the greater good. If you happen to care about the greater good, that's fine. The proposals here work equally well for people motivated by self-interest, national pride, or pure cost-benefit analysis.

This is not a utopian blueprint. It does not describe a perfect world. It describes a world with better coordination infrastructure — which is a much more modest claim, and a much more achievable one.

This is not a tech-solutionist fix. Protocols are tools, not magic. They require human judgment, local knowledge, democratic accountability, and constant maintenance. The metaphors in this document draw from engineering because engineering language is precise and non-ideological — not because the authors believe human coordination reduces to a software problem.

What this *is*: a design pattern. A way of thinking about problems that exceed the scale of existing institutions. A set of observations about coordination mechanisms that already work in specific domains, and a proposal to generalize them. If you prefer: a systems architecture document for a civilization that has outgrown its current operating system but hasn't yet installed the upgrade.

The Prior Question

Before the question of *how* to coordinate, there is a prior question: *what for?*

This document cannot answer that. But it can ask it clearly — and asking it here, before the technical discussion begins, is more honest than tucking it in as a late-arriving philosophical caveat.

Do we — as a species, not as interest groups or nations or generations — want to flourish? Not merely to survive, not merely to grow, not merely to avoid the worst outcomes. To actually flourish: to build a world where the full range of human possibility is available to more than a small fraction of those alive, where the ecological systems that make life possible are stable across the timescales of our children's children, where the intelligence that has emerged in this corner of the universe can keep developing rather than foreclosing itself.

If yes, we need to say so. Collectively. Loudly enough that it shapes what the coordination architecture is built to do.

This is not naive. It is the hardest thing the document asks. The coordination failure tax — the compounding cost of architectures that cannot transmit citizen preferences, that cannot observe slow degradation, that cannot respond at the speed of the threats they face — is also the mechanism by which this question cannot currently be asked at the scale where it matters. We cannot coordinate on what we want because we cannot coordinate. The architecture must be built before the question can be properly answered.

But the question has to come first. Even partially. Even imperfectly. Even in documents like this one, which cannot speak for humanity but can at least name what it hopes humanity might say.

This is an invitation to ask the question — in your community, your institution, your government, your language, your tradition. Not to adopt this document's answer. To contribute yours. The architecture we build will reflect the answers we give. If we do not give any, it will reflect the answers of whoever fills the silence.

The Gap

In 2026, a single shipping container stuck in the wrong canal can cost the global economy billions per day. A virus mutating in one city can shut down schools on the opposite side of the planet within weeks. A financial product designed in London can collapse pension funds in countries whose regulators have never heard of it. A carbon molecule emitted in Houston warms a rice paddy in Bangladesh.

These are not failures of morality. They are failures of architecture.

Our systems of global coordination — the institutions, treaties, and decision-making structures we depend on — were designed for a fundamentally different world. Most of them trace their basic architecture to the Peace of Westphalia in 1648, which established the principle of sovereign nation-states as the primary units

of international order. It was a reasonable design for a world where a message took weeks to cross an ocean and most problems stayed local. The trouble is, the design assumptions no longer hold. Problems have gone global. The coordination infrastructure has not.

This produces a specific, recurring pattern. A crisis emerges that crosses national boundaries — a pandemic, a financial meltdown, a climate threshold, an unregulated technology. No single nation can solve it alone, because the problem operates at a scale that exceeds any national jurisdiction. But no existing international institution has the authority, speed, or trust to coordinate an effective response. So we get a scramble: emergency summits, ad hoc coalitions, blame cycles, and slow-motion failures that everyone can see coming but nobody can prevent. Not because the people involved are incompetent or malicious, but because the coordination architecture wasn't built for this.

Consider the pattern across recent crises. COVID-19 killed millions and cost the global economy an estimated fourteen trillion dollars — not primarily because the virus was unstoppable, but because nations hoarded supplies, withheld data, closed borders unilaterally, and competed for vaccines while the virus mutated freely across all of them. Climate negotiations have produced thirty years of conferences and agreements that have not yet bent the global emissions curve, because the incentive structure rewards defection and punishes early movers. AI is developing faster than any regulatory body can track, in a jurisdictional vacuum where the most consequential technology in human history is governed by a patchwork of national rules that don't talk to each other.

The common thread is not bad leadership. It is structural mismatch. The system is doing exactly what it was designed to do. The problem is that the design specifications are three centuries out of date.

There is a way to see this that bypasses ideology entirely. Some problems are defined by physics rather than politics. The atmosphere does not respect borders. Viruses do not respect borders. Orbital debris does not respect borders. Financial contagion does not respect borders. When the system operates at planetary scale, coordination must exist at planetary scale. This is not a moral claim. It is an engineering observation, as banal and unavoidable as noting that you need a fire department that covers the whole building, not just individual apartments.

So why now? Why does this document exist in 2026 rather than 1996 or 2056?

Several things have converged. The internet has made real-time global coordination technically feasible at near-zero marginal cost — something that was physically impossible for most of human history. AI is increasing the complexity and speed of systemic risks faster than institutions can adapt, making the coordination gap more dangerous with each passing year. Global interdependence is no longer an abstraction; it is a lived experience for anyone who has waited for a delayed shipment, watched a distant war move their

gas prices, or worn a mask because of a virus that originated on another continent. And the cost of coordination failure is now measurable: trillions per crisis, millions of lives, compounding ecological damage.

But there is something the urgency framing usually omits. The failures are not independent. When the financial system's coordination fails, it consumes the fiscal space that health coordination requires. When health coordination fails, it generates the economic disruption that degrades the political stability that climate coordination depends on. When ecological coordination fails, it produces the resource pressures that intensify the conflicts that make every other coordination harder. The crises don't add — they compound. This is not a metaphor. It is the structural mathematics of simultaneous failures in interdependent systems. And it means the case for building better coordination infrastructure is not merely that each individual failure is costly. It is that the interactions between failures are making everything worse faster than any single-domain response can compensate.

The math has changed. The question is no longer whether better coordination is desirable. The question is whether we build it deliberately or continue paying for its absence through catastrophe.

The Observability Problem: Why Institutions Are Structurally Blind

Before describing the solution, it is worth being precise about the problem. The standard account — institutions were designed for a simpler world and haven't kept up — is true as far as it goes. But it doesn't explain why institutions that know they're failing continue to fail. It doesn't explain why organisations staffed by brilliant, committed people consistently miss the crises that retrospective analysis shows were entirely predictable. And it doesn't explain why reform efforts, however well-designed, so frequently fail to produce lasting change.

The deeper answer is observational. Every governance institution reduces the complexity of the environment it governs into a manageable set of signals — the metrics it tracks, the categories it recognises, the data it collects. This compression is not a flaw. It is a necessity. No institution can process everything. But the compression is lossy: the information excluded from the institution's observation channel doesn't cease to exist. It accumulates as an externality, invisible to the institution, until it forces itself into visibility through crisis.

Call this the **observability gap**: the mismatch between what the institution can see and what determines the outcomes it is supposed to govern.

The observability gap is not new. What is new is its scale and speed. For most of institutional history, the gap was manageable because the environment changed slowly enough that crises provided adequate warning before they became catastrophic, and institutions could update their observation channels between crises. The information revolution has compressed this adaptation window dramatically. Problems now move from emergence to crisis faster than institutional reform cycles can complete. The lag that was once manageable has become structurally destabilising.

Three failure modes follow from the observability gap, and all three appear repeatedly across the domains where coordination is failing.

The averaging problem. When institutions aggregate local information into national or global indicators, they see the mean but not the distribution. A national average that looks acceptable can coexist with regional crises in the unmeasured tails. A global indicator that shows improvement can mask the specific communities where conditions are deteriorating fastest. Policies calibrated to the average are systematically miscalibrated to the places where crises concentrate. This is not incompetence. It is the structural consequence of an observation channel designed to see averages.

The speed mismatch. Different problems move at different speeds. Financial contagion moves in hours. Pandemic spread moves in weeks. Ecological degradation moves in decades. Most governance institutions operate at one characteristic speed — the speed of their policy cycle, their reporting requirements, their legislative calendar. Problems that move faster than this speed are invisible until they overwhelm the system. Problems that move slower are visible but politically difficult to sustain attention to across multiple political cycles. The institution is tuned to one frequency in an environment that requires response across many.

The excluded dimension problem. Every institution observes some dimensions of the system it governs and not others. Central banks observe inflation and output while financial fragility accumulates in the unmonitored dimensions. Health ministries track hospital capacity while the social determinants of the next pandemic accumulate invisibly. Regulatory bodies track current AI capabilities while alignment risk accumulates outside their measurement architecture. The crisis, when it comes, occurs along the excluded dimension — which is why it always appears surprising to the institution even when it was predictable to observers using wider observation channels.

Understanding these failure modes changes what better coordination requires. It is not primarily faster institutions, or better-funded institutions, or more committed leadership. It is institutions with observation channels whose dimensionality is adequate to the systems they govern — institutions that can see the distribution, not just the mean; the full frequency spectrum, not just their preferred speed; and the dimensions that matter, not just the dimensions they have historically measured.

Coordination Capacity

If the previous sections described the problem, this one names the solution — or at least the category of solution.

The concept we need is **coordination capacity**: the ability of a system to align diverse, independent actors toward shared outcomes without requiring central command.

This is distinct from several things it might be confused with. It is not *cooperation*, which implies voluntary niceness and falls apart the moment incentives diverge. It is not *governance*, which implies authority — someone governing someone else. And it is emphatically not *world government*, which implies a single centralized body with power over everyone. Coordination capacity is something else: the structural ability of a system to produce aligned behavior from independent agents, each acting in their own context and interests, through shared protocols rather than shared authority.

The most useful analogy is biological. Your nervous system coordinates trillions of cells across dozens of organs without any single organ being in charge. Your brain doesn't micromanage your liver. It doesn't need to. The coordination happens through signals, feedback loops, and shared protocols — the biochemistry of hormones, nerve impulses, immune responses. Each organ does its own work. The nervous system ensures the work is aligned. If you tried to run a human body the way we currently run international relations — each organ sovereign, no shared signaling, ad hoc negotiations every time a crisis hits — you would be dead in minutes.

The technological analogy is equally instructive. The internet runs on a protocol called TCP/IP. Nobody governs the internet centrally. There is no Internet President. No Internet Army. Yet billions of devices coordinate flawlessly every second, routing packets across continents through networks owned by competing companies in rival nations. This coordination emerges not from authority but from protocol design. TCP/IP doesn't care what's in the packets. It doesn't care who sends them. It provides a thin, reliable coordination layer that allows infinite diversity of activity on top of it. The result is the most successful coordination infrastructure in human history — and almost nobody thinks of it as "governance."

Or consider open-source software. Linux runs the majority of the world's servers, phones, and embedded systems. It was built by thousands of contributors across dozens of countries with no central plan, no corporate hierarchy, and no enforcement mechanism beyond shared standards and peer review. The coordination emerges from the protocol: version control, code review, modular architecture, and a culture of transparent iteration. It is not efficient in the way a corporation is efficient. It is resilient in a way no corporation can be.

These analogies point to a principle that is central to everything that follows: **effective planetary coordination and strong local autonomy are not opposites. They are prerequisites for each other.** Without coordination, local autonomy gets crushed by global forces it cannot control — pandemics that overwhelm national health systems, financial contagion that wipes out local economies, climate shifts that make regional agriculture impossible. Without local autonomy, coordination becomes brittle central planning — slow, unresponsive, and disconnected from the reality it claims to manage. The coordination layer and the local layer need each other the way TCP/IP needs the applications that run on top of it.

This leads to a concept worth naming explicitly. We introduced the **Coordination Failure Tax** in the short version — the real economic and human cost of systems that cannot coordinate at the scale of their interdependence. It is worth dwelling on, because it reframes the entire conversation.

The usual framing of global coordination treats it as a cost: something expensive and idealistic that we might undertake if we could afford it. The Coordination Failure Tax inverts this. We are *already* paying for coordination — we are just paying through crisis, redundancy, and collapse instead of through design. COVID's fourteen trillion dollars was a coordination failure tax. Every duplicated research program, every redundant regulatory filing, every supply chain disruption that could have been anticipated with shared data — these are coordination failure taxes. The question is not whether to invest in coordination. The question is whether to invest intelligently or continue paying catastrophically.

And here is the crucial reframe: **the coordination layer is not hypothetical. It already exists in partial form.** Airline safety runs on global coordination protocols — ICAO standards, shared incident databases, universal training requirements — and nobody calls it world government. They call it "the reason your plane doesn't crash." Internet routing, scientific collaboration, financial clearing, pandemic surveillance: these are all domains where protocol-based coordination already operates at global scale, without central authority, and works. What this document proposes is not invention. It is generalization: extending patterns that already work in specific domains to the planetary challenges that currently lack them.

One caveat on the engineering metaphors, because it matters. Protocols for coordinating data packets are simpler than protocols for coordinating human beings. TCP/IP works partly because it doesn't care what's in the packets — it respects the diversity of content by being agnostic to it. Coordination protocols for human systems need an analogous property: they must be flexible enough to respect local context, culture, knowledge, and priorities while enabling alignment on shared challenges. A protocol that works for water management in Sweden may need to look very different in practice from one that works in Rajasthan, even if the underlying principles — shared data, local decisions, consequences for defection — are the same. The protocol is the skeleton. Human judgment and local knowledge are the muscle. Neither works without the other.

This is what we mean by coordination capacity. Not a new authority. Not a new ideology. A new *infrastructure* — one that lets diverse, sovereign actors coordinate effectively on shared challenges while retaining full control over everything else. The question for the rest of this document is: how does that actually work in practice?

How It Actually Works

The easiest way to understand coordination protocols is to look at places where they already exist.

The Montreal Protocol. In the 1980s, scientists discovered that chlorofluorocarbons were destroying the ozone layer — a thin atmospheric shield without which life on Earth's surface becomes untenable. The problem was textbook coordination failure: CFCs were cheap, useful, and produced by companies in many countries. No single nation banning them would solve the problem, and any nation that banned them unilaterally would put its industries at a competitive disadvantage. Classic prisoner's dilemma.

What happened next is instructive. The Montreal Protocol, signed in 1987, established shared monitoring of ozone depletion (so everyone was looking at the same data), set phase-out timelines that nations implemented according to their own industrial contexts, and created trade restrictions that made it more expensive to keep producing CFCs than to switch alternatives. No world government was created. No sovereignty was surrendered. Within two decades, CFC production dropped by over 99%, and the ozone layer began to recover. It is arguably the most successful piece of global coordination in human history, and most people have never thought about it.

Global airline safety. Every day, over a hundred thousand commercial flights operate worldwide with an astonishingly low fatality rate. This is not because pilots are heroes (though some are). It is because the International Civil Aviation Organization maintains a set of coordination protocols — shared safety standards, mandatory incident reporting databases, universal training requirements, standardized communication procedures — that every airline in the world follows. When a near-miss happens in Brazil, the data enters a global database that an engineer in Japan can learn from. Nobody thinks of ICAO as "global governance." It is simply engineering that keeps people alive.

The internet's root DNS system. The Domain Name System that translates human-readable web addresses into machine-readable IP addresses is managed by thirteen root server clusters operated by independent organizations across multiple countries. There is no Internet President. No single entity controls the system. Yet it handles trillions of queries daily with extraordinary reliability. The coordination emerges from protocol design — shared standards, distributed redundancy, transparent governance — not from central authority.

CERN. Twenty-three member states fund and operate the largest particle physics laboratory on Earth. Thousands of scientists from competing nations — including nations with active geopolitical tensions — collaborate at extraordinary trust levels on experiments that cost billions. The coordination works because the protocols are clear, the data is shared, and the incentive structure rewards participation: if you want access to the most advanced physics infrastructure in the world, you contribute to it. High-energy physics doesn't care about your flag.

WHO's Global Influenza Surveillance and Response System. For decades, national laboratories have voluntarily shared influenza virus samples and genomic data through a coordinated network. The shared visibility allows early detection of novel strains. The local execution means each nation runs its own surveillance and response. The benefit of participation — access to shared data, early warning, vaccine development inputs — outweighs the cost of contributing. It is imperfect and has well-documented limitations, but it works well enough that when it fails (as parts of it did during COVID, when data sharing broke down), the consequences are measured in millions of lives.

These examples come predominantly from Northern-led institutions, and that is a bias worth noting. Coordination traditions exist far beyond the Western institutional canon. ASEAN's consensus-based decision-making has maintained regional stability across extraordinary cultural and political diversity for decades. Indigenous resource management systems — from Pacific Islander fishing protocols to Aboriginal Australian fire management — sustained coordination across communities for millennia before Westphalia was conceived. The African Union's continental frameworks for health coordination and conflict mediation operate in contexts of complexity that European institutions rarely face. A complete account of "how coordination works" would draw deeply from these traditions. This document can only gesture toward them and invite their practitioners to contribute to what comes next.

What matters here is the pattern. Strip away the specifics of ozone chemistry or DNS architecture, and every successful coordination system shares three structural properties.

First: shared visibility. All relevant actors can see the same data. The ozone measurements. The genomic sequences. The flight incident reports. The particle collision results. Shared visibility doesn't mean everyone agrees on what the data means — it means everyone is arguing from the same facts rather than from competing narratives. This alone eliminates a huge category of coordination failure. Most international disputes are not actually about values. They are about information asymmetry.

Second: local execution. No central authority tells anyone what to do. Local actors — nations, communities, organizations, individuals — execute based on their own context, priorities, and incentives. The Montreal Protocol didn't prescribe how each country should phase out CFCs. It set targets and let each nation find its own path. ICAO doesn't fly the planes. Airlines fly the planes, using shared standards. This is critical: it means the coordination layer doesn't need to understand every local context. It only needs to make the shared information visible and the incentive structure clear. Local knowledge handles the rest.

Third: meaningful cost of defection. This is the one that makes the difference between a protocol and a suggestion. In every successful example, ignoring the coordination is more expensive than participating in it. CFC producers that kept manufacturing faced trade restrictions. Airlines that hide incidents lose insurance coverage and landing rights. Nations that don't contribute influenza data lose access to shared surveillance. CERN members that don't pay their share lose access to the facilities. The cost of defection doesn't have to be punitive. It just has to be real.

An honest document has to address an obvious objection here: these examples worked partly because of underlying power dynamics that the tidy three-part pattern doesn't capture. The Montreal Protocol succeeded in part because the United States — the world's largest CFC producer — decided it was in its economic interest to switch to alternatives its companies were already developing, and then used trade leverage to bring others along. ICAO standards work because airlines that ignore them can't get insurance, and insurers operate in markets shaped by powerful states. Protocols without consequences are suggestions. And the design of enforcement mechanisms — who decides what counts as defection, who imposes costs, who adjudicates disputes — is where most coordination attempts either fail or get captured by powerful actors.

This document does not pretend this problem is solved. It is the hardest problem in coordination design, and the specific mechanisms matter enormously. What can be said is that the answer lies somewhere between "UN strongly-worded letters" (consequences too weak to change behavior) and "world police force" (consequences too centralized to be legitimate). The working models — binding arbitration with real economic stakes, graduated sanctions, transparent accountability mechanisms, loss of access to shared benefits — offer a direction. But the honest truth is that enforcement design for planetary-scale coordination is an active engineering problem, not a completed blueprint. Anyone who tells you otherwise is selling something.

With the pattern established, here is how it might apply in practice.

Water. Two communities share a river. The upstream community's agricultural runoff affects the downstream community's drinking water. Under the current system, this is a political fight — lawsuits, lobbying, blame. Under a coordination protocol: shared sensors monitor water quality in real time (shared visibility). Each community manages its own land use and water treatment (local execution). But the upstream community's water discharge must meet agreed standards, verified by shared monitoring, with graduated consequences for violation — reduced water allocation, financial penalties, loss of standing in regional water governance (cost of defection). The dispute doesn't disappear. But it moves from a political arena, where power determines outcomes, to a protocol arena, where data and pre-agreed rules determine outcomes.

Pathogens. A novel virus appears. Under current systems, the discovering nation has perverse incentives: reporting early brings economic punishment (travel bans, trade restrictions) while concealing buys time. Under a coordination protocol: genomic surveillance networks flag anomalies automatically (shared visibility). Local health systems activate pre-practiced response plans adapted to their own infrastructure and

culture (local execution). Nations that report early receive automatic resource support — medical supplies, technical assistance, economic stabilization — while nations that conceal face exclusion from shared surveillance data and pandemic preparedness benefits (cost of defection). The incentive structure flips: reporting becomes the rational move.

Supply chains. A critical component shortage threatens manufacturing across multiple countries. Under current systems: panic hoarding, export restrictions, bidding wars that drive up prices and leave the poorest nations empty-handed. Under a coordination protocol: shared inventory databases and demand forecasting provide real-time visibility into where shortages are developing (shared visibility). Regional manufacturing hubs adjust production based on local capacity and demand (local execution). Nations that impose unilateral export restrictions during a declared shortage lose preferential access to shared supply chain data and coordination benefits during future disruptions (cost of defection).

In each case, the pattern is the same: make the right information visible, let local actors decide, and ensure that defection costs more than cooperation. That is the core mechanic. The rest — and it is a lot of rest — is implementation detail.

But there is a principle that governs all the implementation detail, and it is simple enough to state in one sentence: **decisions should live at the lowest level capable of handling them.** Your community manages your school. Your city manages your water system. Your region manages your watershed. Your nation manages your laws and infrastructure. Continental bodies manage trade corridors and shared ecosystems. The global coordination layer only touches what physics and interconnection force it to touch: atmosphere, pandemics, financial contagion, orbital debris, and the handful of other problems that genuinely cannot be solved at any lower level. Everything else stays local. Not as a concession. As a design requirement. Because local actors have local knowledge, and coordination without local knowledge is just bureaucracy with a longer reach.

The Observation Trap: Why We Cannot "IPCC" Everything

To understand why the global layer must be restricted this way, consider the difference between governing the climate and governing inequality.

When the international community realized the scale of the climate crisis, it created the Intergovernmental Panel on Climate Change (IPCC). It is widely considered a success of global coordination. Because of this success, policymakers frequently suggest creating "an IPCC for inequality," "an IPCC for AI," or "an IPCC for global health."

This instinct is understandable but reveals a structural mismatch.

The IPCC works because the Earth's atmosphere has a relatively tractable observation problem. The atmosphere is a single, well-mixed fluid system. The key variables (parts per million of CO₂, global mean temperature) can be measured from relatively few vantage points, and a ton of carbon emitted in Beijing behaves identically to a ton emitted in Berlin. The system has low variance.

Human inequality, poverty, and social stability do not have this property. They are hyper-localized, high-variance phenomena. The conditions that produce extreme poverty in a Congolese mining region, a Brazilian favela, and a post-industrial Welsh town are structurally entirely different. They have different historical causes, different cultural feedback loops, and require radically different intervention leverage points.

When a centralized permanent body attempts to govern a high-variance system, it faces immense computational pressure to aggregate these diverse realities into a single global picture.

This is the **Averaging Problem** — the same observability failure described earlier, operating at institutional scale. Aggregation systematically destroys the local signal that makes effective intervention possible. A centrally aggregated inequality index or health policy observes the *average*; it cannot observe the *specific*. Policies calibrated to the average will be systematically mismatched to the conditions that deviate from it — which is precisely where crises concentrate.

This is not a political objection to global unity; it is a structural limitation defined by Ashby's Law of Requisite Variety. A governing body can only respond to the variety of disturbances it is capable of observing. If the global body attempts to govern interventions directly, rather than simply setting boundaries and allowing local nodes (subsidiarity) to execute the response, it will replicate the exact structural failure modes it was designed to solve.

The implication is architectural: the global coordination layer must be thin. Its job is not to govern everything that matters globally. It is to govern the narrow set of shared physical systems where collective action is the only viable response, and to provide the information infrastructure that enables better-matched governance at every other level. The temptation to expand the global layer — to centralize more, observe more, direct more — is the temptation that most coordination failures eventually succumb to. Resisting it is a design discipline, not a political concession.

What Changes, What Doesn't

If the previous sections described what coordination capacity is and how it works, this one addresses the question that most readers actually care about: what does it mean for me?

Start with what doesn't change, because the list is long and reassuring.

You are still Swedish, or Japanese, or Texan, or Kenyan. Your nation still governs itself. Your community still makes its own decisions about education, culture, land use, public life. Your elections still matter. Your traditions are still yours. Local governance stays local — not as a consolation prize, but because the entire architecture depends on it. A coordination system that overrides local knowledge defeats its own purpose.

Cultural diversity is not merely tolerated by this model. It is required. The coordination layer needs cognitive diversity for the same reason ecosystems need biodiversity: monocultures are fragile. A planet where everyone thinks the same way, organizes the same way, and values the same things would be catastrophically vulnerable to any threat its single perspective couldn't anticipate. The protocols described here are designed to be culture-agnostic at the infrastructure level precisely so they can support cultural richness at the human level. TCP/IP doesn't care what language your website is in. Coordination protocols shouldn't care what language your parliament speaks.

What does change is more specific, and it is worth being honest about rather than burying it in reassurance.

You lose the ability to externalize costs. This is the real trade, and it deserves plain language. Right now, a nation can pollute across borders, dump carbon into a shared atmosphere, destabilize a neighboring region and leave others to manage the refugees, or design financial products that generate profits domestically and crises internationally. Coordination protocols don't prevent you from doing things. They prevent you from escaping the consequences. For the vast majority of people and nations — those who are currently on the receiving end of other people's externalized costs — this is straightforwardly good. For those whose current prosperity depends on externalizing damage — certain extractive industries, certain financial models, certain geopolitical strategies — it represents a genuine loss. This document will not pretend otherwise. The question is whether a system that lets powerful actors externalize costs onto everyone else is one worth preserving. We think the answer is obvious, but we recognize that the people benefiting from the current arrangement may disagree.

You lose the ability to free-ride. Shared resources — atmosphere, oceans, financial stability, pandemic preparedness — currently operate on a system where everyone benefits and nobody is required to contribute proportionally. Coordination protocols create accountability for proportional contribution. If you use the shared infrastructure, you help maintain it. This is not radical. It is how every functional commons in human history has worked, from medieval grazing lands to modern condominium associations. What's new is applying the principle at planetary scale.

You lose the ability to selectively defect. Under current arrangements, a nation can participate in coordination when convenient and exit when it isn't, with minimal consequences. Coordination protocols with meaningful cost of defection change this calculus. You can still exit — genuine exit mechanisms are a design requirement, not an afterthought — but exit has costs, just as leaving a trade alliance or an open-

source consortium has costs. The value proposition is that the protocols are better than going it alone. Nobody forces you to participate. But the benefits of participation are real enough that choosing isolation becomes an expensive decision.

Now for the concerns that reasonable people raise about all of this.

Won't this stifle innovation? This is a legitimate worry. Coordination that requires consensus from every actor before anyone can move is a recipe for stagnation. But coordination is not consensus. Good protocol design includes explicit space for experimentation. The internet didn't succeed by coordinating everything — it succeeded by coordinating the minimum necessary infrastructure (packet routing, address resolution) and leaving everything else radically free. The result was the most explosive period of innovation in human history, *because* the coordination layer was reliable. The same principle applies here: coordinate at the infrastructure level, compete and innovate at the application level. Innovation sandboxes — zones where new approaches can be tested without requiring planetary permission — are a feature of the architecture, not an exception to it.

Won't this just be captured by powerful actors? This is more than legitimate — it is the central design challenge. Coordination systems must account for the fact that power is unevenly distributed, and that powerful actors have strong incentives to capture or corrupt coordination mechanisms to serve their own interests. Every international institution in history has faced this problem, and many have succumbed to it. Protocol design cannot eliminate power asymmetry. What it can do is constrain destructive expressions of power through transparency, distributed governance, and structural checks that make capture difficult and visible. Any coordination architecture that ignores the capture problem is naive. This document takes it seriously as a design constraint — not a problem to be solved once and forgotten, but a pressure that requires continuous structural resistance.

Won't this lead to technocratic rule by unelected experts? Another legitimate concern, and one with historical precedent. The European Union's legitimacy crisis offers a cautionary tale: coordination that outpaces democratic consent generates backlash that can destroy the coordination itself. The answer is not less coordination, but coordination with built-in democratic accountability. Concrete mechanisms: sunset clauses requiring periodic renewal by democratic bodies, so no protocol persists without active consent. Transparent governance structures with elected representatives, not just appointed experts. Genuine exit mechanisms that allow nations and communities to leave protocols without catastrophic penalty. If a coordination system cannot survive democratic scrutiny, it should not exist. This is not a concession to populist pressure. It is a design principle. Protocols that lack democratic legitimacy are protocols that will fail — not because they're wrong, but because they'll be rejected by the people they claim to serve.

What about nations that simply refuse to participate? They are free to do so. The opt-in structure means that no nation is forced into any protocol. But the architecture is designed so that participation is more attractive than isolation. A nation that opts out of pandemic surveillance loses access to shared early warning

data. A nation that rejects climate coordination bears the full cost of climate impacts without access to shared adaptation resources. A nation that ignores financial transparency protocols finds its institutions excluded from the networks where capital flows most efficiently. This isn't punishment. It is the natural consequence of choosing not to participate in a system whose benefits flow from participation. The same logic applies to any club, any alliance, any open-source project. The door is always open. But you don't get the benefits from the hallway.

Snapshots

Abstractions are forgettable. What follows are three sketches of the same world seen from different angles — not utopian projections, but plausible Tuesdays. The goal is not to inspire. It is to make the concrete implications of coordination visible in ordinary life.

A Tuesday in 2025: The Path Not Taken

A novel respiratory virus emerges in a densely populated region of Southeast Asia. The local hospital flags an unusual cluster of pneumonia cases, but the national government hesitates to report internationally. The last time a country in the region issued an early warning, it was rewarded with travel bans, trade restrictions, and a cratered tourism sector. The incentive is to wait, assess, and hope it resolves on its own.

It doesn't. Within three weeks, the virus is in twelve countries. By the time the WHO convenes an emergency session, it is in thirty. Nations close borders unilaterally, using different criteria and different timelines, creating chaos at ports and airports. Medical supply shipments are seized by transit countries. Pharmaceutical companies begin parallel vaccine development programs, competing rather than collaborating, each protecting proprietary data. Conspiracy theories about the virus's origin spread faster than the epidemiology. Cable news runs twenty-four-hour blame cycles. Social media platforms amplify fear.

Six months later, the virus is in 120 countries. Schools are closed on four continents. The global economy has lost trillions. Retrospective analyses will show that early, coordinated action in the first two weeks could have contained it regionally. But early coordinated action required infrastructure that did not exist: shared surveillance, pre-agreed response protocols, incentives for early reporting, and automatic resource support for affected regions. Nobody built it. Everybody paid for its absence.

This is not a hypothetical. Versions of this story have played out repeatedly — with SARS, H1N1, Ebola, and COVID-19. The details vary. The structural failure is identical every time.

A Tuesday in 2045: The Coordinated Response

The same virus emerges. The same hospital flags the same unusual cluster.

This time, the genomic sequence is uploaded within hours to a shared surveillance network. Not because the local doctors are more virtuous, but because the protocol rewards early reporting: nations that flag anomalies early receive automatic resource support — medical supplies, technical assistance, economic stabilization funds — while nations that conceal data lose access to shared surveillance benefits. The incentive structure has been flipped. Early warning is the rational move.

Within seventy-two hours, the sequence is being analyzed in laboratories across four continents. Pre-agreed response protocols activate at regional level. Travel advisories are data-driven and proportionate — targeted screening rather than blanket bans. Affected communities receive support; they are not punished for transparency. Local health systems execute response plans they have already practiced, adapted to their own infrastructure and cultural context.

The virus is contained regionally within six weeks. Schools stay open in Stockholm and São Paulo and Nairobi. Financial markets dip two percent and recover within a month. The WHO publishes a routine after-action report. Nobody writes a book about it. Nobody makes a documentary. Six months later, most people have forgotten it happened.

Not because anything heroic occurred. Because the protocols worked. Because the infrastructure existed. Because someone, years earlier, decided to build it.

A Tuesday in 2045: The Quiet Ordinary

In Telangana, a farmer named Priya checks her phone over morning tea. A shared satellite data feed, accessible through an open protocol, shows shifting rainfall patterns for the coming weeks. She adjusts her planting schedule accordingly — something her grandmother did by intuition, but with less precision and more risk. The data comes from a constellation of weather satellites funded by a global commons mechanism. Priya doesn't think about this. She thinks about her crops.

In Curitiba, a city council reviews its monthly watershed report. Three neighboring municipalities share the same river basin, and an open-source coordination tool — developed collaboratively by engineers in South Korea and Kenya and adapted locally — lets them manage water allocation in real time. There was a dispute last year about upstream irrigation rights. The shared monitoring data resolved it in a week. Under the old system, it would have been a lawsuit lasting years.

In Mombasa, a seventeen-year-old named Amara contributes a data set on drought-resistant millet varieties to an open agricultural research commons. Her school's contribution is logged, and her community earns recognition in a knowledge-sharing network that gives participating communities early access to new crop

research. Amara wants to study agricultural engineering. She doesn't think of herself as participating in "global governance." She is doing a school project that happens to use tools that work.

None of these scenes are dramatic. That is the point. The best possible future is not one of grand heroism and world-historical speeches. It is one where the infrastructure is so well-designed that it becomes invisible — where coordination is as unremarkable as plumbing, and as essential. A future that is, in the best sense of the word, boring.

The Honest Part

If you have read this far, you deserve some candor.

You might be thinking that this sounds impossibly ambitious. That redesigning planetary coordination is the kind of thing that gets discussed at Davos and then nothing happens. That this is a nice thought experiment with no realistic path to implementation. That the gap between "a protocol document written by some guy" and "actual functioning global infrastructure" is so vast as to make the exercise pointless.

You would not be entirely wrong. By the standards of how institutional change usually happens — slowly, through existing power structures, driven by crisis rather than foresight — the ideas in this document are impractically ambitious. Most readers will put this down and forget about it by Thursday. That is fine. This document is not for most readers. It is for the small fraction who read the preceding sections and felt something click — a recognition that the problems described are real, that the patterns identified are accurate, and that the direction proposed is at least worth thinking about seriously. If you are in that fraction, this section is for you. If you are not, thank you sincerely for reading this far.

Let us be honest about who wrote this. This document originated with a systems thinker in Sweden who spends his evenings playing strategy games on Discord and his days worrying about global governance design. He is not the UN Secretary-General. He does not run a think tank or a foundation. He does not have a policy staff, a research budget, or a Rolodex of world leaders. He has a desktop computer, access to five large language models with more patience than most human collaborators, and the conviction that someone should at least write this down — even if the "someone" turns out to be some guy in Upplands Väsby.

This is mentioned not for self-deprecation, but because it is relevant to the argument. If coordination capacity matters, then it should not require elite credentials to describe it. If the ideas here are sound, they should be sound regardless of who articulated them. And if they are unsound, the author's lack of institutional prestige is the least of the problems. Judge the argument, not the résumé.

Now for the hard admissions.

This document is probably wrong in important ways. The map is not the territory. The analogies are imperfect — human coordination is messier than packet routing, and the differences matter. The case studies are cherry-picked in the sense that all case studies are: we chose examples that illustrate the pattern, not examples that contradict it. The proposed mechanisms are sketched, not engineered. The enforcement problem remains genuinely unsolved. If you find errors, gaps, or fatal flaws, that is not a failure of the document. It is the document working as intended. The protocol for improving it is the same as the protocol it describes: diverse perspectives, honest feedback, and shared commitment to getting closer to something that works. Bring your map. We'll compare.

The critics have real points. The game theorist is right that protocols without enforcement mechanisms are suggestions, and that in high-stakes domains rational actors will defect when it serves their immediate survival. The sovereignty defender is right that coordination can become technocratic capture, and that "neutral" protocols have a way of serving the interests of whoever designed them. The post-colonial critic is right that a document written from a Northern European perspective using Western-trained AI models embeds biases it cannot fully see, and that the history of "universal" frameworks being imposed on the Global South is long and ugly. The accelerationist is right that over-coordination stifles experimentation and that some of history's greatest advances came from actors who broke protocols rather than followed them.

This document does not have complete answers to any of these critiques. What it has is a framework that takes them seriously as design constraints rather than dismissing them as obstacles. Enforcement must be graduated, legitimate, and transparent. Democratic accountability must be structural, not decorative. Cultural and positional biases must be named and actively countered. Space for experimentation must be built in, not bolted on. These are engineering requirements, not political concessions. Whether this document meets them adequately is for readers — especially readers whose perspectives differ most from the author's — to judge.

The bootstrapping problem is real. You need coordination to build coordination tools. You need trust to build trust infrastructure. You need a critical mass of participants before the benefits of participation outweigh the costs. This is the "first penguin" problem — someone has to jump in before anyone can see whether the water is safe. We do not have a clean answer to this. What we have is a historical observation: every coordination system that now seems inevitable — the internet, airline safety, international mail — started with a small group of actors who decided to try it and see. ARPANET had four nodes. The Universal Postal Union started with twenty-two countries. The Montreal Protocol was signed by twenty-four nations, most of whom were hedging their bets. The pattern is not "universal agreement followed by implementation." The pattern is "a small committed group builds something that works, and others join because the benefits become undeniable." That is the model here. Not mass adoption. Not top-down imposition. A seed crystal.

Why we think the math might be different this time. Not certainly different. Possibly different. The internet has reduced the cost of coordination by orders of magnitude. Information that once took months to share now takes milliseconds. Tools for transparent monitoring, shared data, and distributed decision-making exist at a scale and sophistication that previous generations could not have imagined. The cost of

coordination failure is now so high — trillions per crisis, accelerating ecological damage, AI risks that compound annually — that the case for inaction is weaker than it has ever been. And a generation is coming of age that thinks natively in networks, open-source collaboration, and distributed systems, not in hierarchies and territorial boundaries. None of this guarantees success. But it changes the odds enough to make the attempt worth making.

The Transition Mechanics: Upgrading Legacy Hardware

When the structural failures of centralised governance become undeniable, a dangerous false dichotomy usually emerges: we must either fiercely defend the failing status quo, or burn the institutions to the ground and start over.

Both paths are catastrophic. Defending a centralised bottleneck guarantees slow systemic collapse driven by latency and the averaging problem. Destroying the institutions — the revolutionary impulse — eliminates centuries of accumulated social capital, physical infrastructure, logistical networks, and institutional memory. When complex systems lose their legacy infrastructure, they do not evolve; they regress into power vacuums that are immediately filled by authoritarianism.

Institutions like the United Nations, the nation-state, and municipal councils are not the enemy. They are the **legacy hardware** of human coordination. The problem is not that the hardware exists; the problem is that it is running an outdated, high-latency operating system. We do not need to destroy the computer. We need to install better routing protocols.

The transition to a more distributed, polycentric governance model happens through a process of **structural permeation**, enabled by a specific upgrade to our baseline economic infrastructure.

The immune system problem

Here the document needs to be honest about something that earlier versions skated past.

The institutions defending the existing architecture are not simply failing. From the perspective of those they serve, many of them are working precisely as intended. The actors who profit from cost externalisation — extractive industries, certain financial models, geopolitical strategies that depend on asymmetric power — have strong interests in the continuation of the current arrangement. And those actors have disproportionate influence over the institutional forms through which coordination proposals are received and processed.

This is not a conspiracy. It is the predictable output of an incentive architecture that rewards those who benefit from the status quo with the resources and influence to defend it. Every coordination proposal enters an institutional environment shaped by those interests. The ones that threaten the most profitable externalisations face the most determined resistance. And the resistance doesn't usually take the form of direct opposition — it takes the form of absorption. The proposal is welcomed, studied, amended, piloted in a constrained form, and gradually defused until it no longer threatens the underlying arrangement. The language of reform is adopted; the structure of reform is not.

A transition strategy that doesn't account for this will be absorbed by the same machinery it is trying to change. The design requirement is not merely to build better protocols. It is to build them in ways that are structurally difficult to absorb — that create constituencies for the new architecture, that make the dysfunction of the old architecture increasingly visible and costly, and that expand the political coalition for structural change faster than the immune response can neutralise it.

Concretely, this means three things. First, demonstrate value at the periphery before demanding adoption at the centre. Pilot protocols in willing municipalities, willing industries, willing nations — and design the pilots explicitly to generate evidence that the surrounding architecture cannot easily dismiss. Second, build bypass architectures with sunset conditions. A bypass that works too well relieves pressure on the unreformed core, allowing the core to survive longer in its dysfunctional state. Good bypass design increases pressure on the core rather than relieving it: the municipal laboratory that produces better outcomes makes the national system's dysfunction more visible, not less. Third, expand the observation channel that the political system uses to evaluate institutional performance. The immune system's most powerful defence is making governance failure invisible — through aggregated statistics that hide distributional consequences, through indicators calibrated to dimensions on which incumbent actors perform well, through attribution chains that are too long to trace accountability. Shared visibility — the first pillar of coordination capacity — is also the first requirement for building political demand for structural change.

Universal basic income as cognitive infrastructure

In a traditional centralised institution, the human nodes — civil servants, administrators, citizens — are caught in a survival dynamic. Because their livelihood depends on the institution, their primary incentive is compliance, not adaptation. A node cannot afford to challenge a failing system if doing so threatens their baseline survival. The cognitive bandwidth of the system is consumed by institutional self-preservation rather than problem-solving.

To upgrade the legacy hardware, we must alter the physics of this constraint. This is one structural purpose of a universal basic income.

A guaranteed income floor is not merely a mechanism for poverty alleviation; it is foundational cognitive infrastructure. By decoupling baseline survival from institutional obedience, it lowers the noise floor of survival anxiety — freeing the intellectual bandwidth that nodes need to actually engage with complex problems rather than manage personal risk. More precisely: the rational behaviour of institutional actors who cannot afford to challenge the system is the same rational behaviour that makes immune systems so effective. Reducing the personal cost of dissent changes the incentive architecture that makes immune system behaviour rational. At scale, this changes what institutions are capable of perceiving and acknowledging about themselves.

The permeation strategy

Individuals with that security act as **adaptive nodes** within rigid legacy structures. Rather than fighting institutions from the outside, they permeate them.

Equipped with diagnostic thinking about where information actually lives, where decision latency is highest, and where the averaging problem is generating the most damage, these nodes begin rewiring the circuitry while the machine is still running:

1. **Lateral routing:** Building horizontal communication channels that bypass slow vertical hierarchies, allowing municipalities and local communities to coordinate directly across jurisdictions.
2. **High-fidelity feedback loops:** Implementing localised sensing and response mechanisms that move decision authority to the level where the relevant information exists — neutralising the averaging problem from within.
3. **Resource redirection:** Using the institution's existing reach — its budgets, legal frameworks, physical infrastructure — and orienting it toward outcomes that the institution's formal mandate already supports, but that its architecture currently prevents it from achieving.

This is how systemic transformation actually occurs. Not through replacement, but through inside-out upgrade. The reach, scale, and stability of legacy institutions are preserved; their centralised bottlenecks are gradually redistributed toward faster, higher-fidelity distributed protocols.

We do not need a revolution to navigate the coordination crisis. We need a coordinated firmware update, executed by sovereign nodes from within — designed with explicit attention to the immune systems that will resist it, and structured to build political demand for structural change rather than simply to demonstrate that change is possible.

An Invitation

This document does not end with a call to action. There is no petition to sign, no organization to join, no movement to enlist in.

It ends with an invitation. Several, actually, depending on who you are.

If you are a **thinker** — an academic, a researcher, a policy analyst, someone who enjoys taking ideas apart to see how they work — the invitation is to critique this. Rigorously. Find the load-bearing assumptions and test them. Steel-man the strongest objections and see if the framework survives. Identify what's missing, what's naive, and what's wrong. The most valuable thing you can do is tell us where this breaks. And if you want a specific target: find a domain where the three-part pattern — shared visibility, local execution, cost of defection — has been fully implemented and has failed to produce coordination. That is the falsifying evidence this framework needs to encounter.

If you are a **builder** — an engineer, a designer, an organizer, someone who makes things work in practice — the invitation is to prototype. Pick a coordination failure in your domain — your city's water management, your industry's supply chain, your community's disaster preparedness — and design a protocol for it. Use the pattern: shared visibility, local execution, cost of defection. See if it works. Report back what you learn. And when you build a bypass around a dysfunctional core, design it with intention: not just to work, but to make the core's dysfunction undeniable.

If you are a **sharer** — someone with a platform, a community, a network, or just a group chat — the invitation is to translate and adapt. This document was written in English from a European perspective. It needs to be rewritten by people in Jakarta and Accra and Medellín and Mumbai. Not translated literally — *adapted*, so the examples are local, the concerns are real, and the voice is authentic. Fork it. The source is open. The license is permissive. The document's own argument demands that it evolve beyond its origins.

And if you are none of these — if you are simply someone who read this and found it interesting — the invitation is smaller but no less important. Notice coordination failures when you encounter them. The next time an institution fails in a way that feels predictable and preventable, ask: what protocol would make this unnecessary? What information would need to be shared? Who would need to act locally? What would make defection costly? And ask the harder question too: whose interests are served by the current arrangement, and what would it take for them to find the new arrangement more attractive than the old one? You may find that thinking in these terms changes what you see.

For those interested in the technical depth behind these ideas — specific protocol designs, economic mechanisms, governance architectures, and implementation roadmaps — a comprehensive open-source framework ecosystem is in development. It is called the Global Governance Frameworks project, and it is

available at [globalgovernanceframeworks.org](https://www.globalgovernanceframeworks.org) (<https://www.globalgovernanceframeworks.org>). It contains detailed specifications for many of the mechanisms sketched in this document: coordination councils, enforcement architectures, economic transition protocols, democratic accountability structures, and much more. It is a work in progress, built in the open, with all the rough edges that implies.

It is, like this document, a starting point. Not a destination.

Further Reading

For those who want to go deeper, a few starting points — not an exhaustive bibliography, but a set of doors.

Elinor Ostrom — *Governing the Commons* (1990). The foundational work demonstrating that humans can and do coordinate shared resources without either central authority or privatization. Ostrom documented hundreds of real-world commons — fisheries, irrigation systems, forests — governed successfully by local protocols. She won the Nobel Prize in Economics for proving that the "tragedy of the commons" is a design failure, not an inevitability.

The RFC Series (1969–present). The Request for Comments documents that built the internet. Not a book but a living tradition: thousands of technical specifications, written in plain language, openly debated, iteratively improved, and freely available. If you want to understand what protocol-based coordination looks like in practice — the culture of it, the tone, the method — read a few RFCs. Start with RFC 1, written by Steve Croker in 1969, which begins with a note about the "hope that the mere act of writing down our discussions will encourage others to do the same."

The Federalist Papers (1787–1788). Eighty-five essays arguing for the ratification of the United States Constitution, written by Alexander Hamilton, James Madison, and John Jay under a shared pseudonym. Relevant not because the American system is a model to copy, but because the papers represent one of history's most successful examples of persuasive writing about coordination architecture. They had to convince skeptical, sovereign states to voluntarily adopt shared protocols. The rhetorical challenges they faced are remarkably similar to the ones this document faces. They also had a print run of about 500 copies.

Stafford Beer — *Brain of the Firm* (1972). A cybernetics classic that models organizations as viable systems requiring specific coordination structures to remain adaptive. Beer's Viable System Model — which describes how any complex system needs recursive layers of coordination, each with appropriate autonomy — is one of the intellectual foundations for the ideas in this document, whether or not the reader has encountered it before.

Donella Meadows — *Thinking in Systems: A Primer* (2008). The most accessible introduction to systems thinking available. Meadows writes with clarity and warmth about feedback loops, leverage points, and the behavior of complex systems. If the ideas in this document interest you but the engineering language feels unfamiliar, start here.

Ross Ashby — *Introduction to Cybernetics* (1956). The formal source of the Law of Requisite Variety that underlies the observability argument in this document. Harder going than Meadows, but the chapters on variety and regulation are worth the effort for anyone who wants the mathematical foundations of why coordination requires matched complexity between the governing system and the governed one.

This document is version 0.3, released under Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0). You are free to share, adapt, translate, critique, fork, and improve it, provided you give appropriate credit and distribute your contributions under the same license. The source is available as a Markdown file at <https://github.com/BjornKennethHolmstrom/website/tree/main/src/lib/content/the-coordination-imperative-v0.3.md> (<https://github.com/BjornKennethHolmstrom/website/tree/main/src/lib/content/the-coordination-imperative-v0.3.md>).

If five AIs and a Swede can get this far, imagine what happens when the rest of you join in.