

Sovereignty Without Dominance

Rightful authority at human scale, in the age of AI

Précis. This is an essay about a single word. *Sovereignty* has been quietly conscripted into the contest between great powers, where it has come to mean the capacity to build, own, and out-compete — the power to dominate a technology. There is a better way to hold the word, once it is set beside *legitimacy* — the rightful authority it is so often asked to carry. The sovereignty worth wanting then becomes plain: a community’s control over the data and systems that act on it, held with legitimacy, at human scale, federated rather than raced, and never surrendered. The argument extends Alvin Wang Graylin rather than rebutting him — a different lane, not a contrary one; he is right that the race is real and dangerous — and it asks Western readers to do something they have been trained against: to turn the same sceptical gaze they reserve for China onto the United States, and onto the structure of dependency itself, with respect for every party and enmity toward none. It closes with a live exhibit, drawn from the days in which it was written.

The argument everyone is having

For two years the public conversation about artificial intelligence has been organised as a race between two powers. A small nation, a regional bloc, a community is told its choices reduce to a question of allegiance: the American stack or the Chinese stack. Build on one, and accept its terms; build on the other, and accept its terms instead.

Alvin Wang Graylin — a technologist with three decades of operating experience across the United States and China — has made the most lucid case against this framing, and he is right. Writing with Paul Triolo, he argues that “the AI competition is increasingly being framed within narrow national security terms, as a zero-sum game,” and that “there will not and cannot be any long-term winners if the intense competition continues on its current path.” His correction is precise: the obsession is with the wrong verb. “In this race,” he writes, “adoption beats invention. Distribution beats dominance.” A global effort “akin to the CERN for AI,” he proposes, “will bring much more value to the world, and a peaceful end, than a Manhattan Project for AI.”

Notice the word Graylin attacks. It is not *sovereignty*. It is *dominance* — the pursuit of an unassailable, permanent lead. He is careful, too, about what bifurcating the world costs: splitting the AI ecosystem into rival spheres, he warns, “will only create more safe spaces for bad actors to hide and strike.” This is the considered view of someone who knows both systems from the inside, and the Village project begins by agreeing with all of it. The mindset of scarcity and dominance is the obstacle, not the technology.

I should be careful, though, not to mistake a quarrel with the race for a denial of it. The race is not a misperception to be corrected; it is an accurate description of *what is emerging* — Graylin and Triolo wrote in part to raise the alarm about exactly that — and the months since have only sharpened it. In March 2026, Iranian drones struck Amazon’s data centres in the United Arab Emirates and Bahrain, the first known military strikes on an American hyperscaler’s infrastructure, with critical data infrastructure now, in one analyst’s phrase, “emerging as a potential target in modern warfare.” In June, the European Union tabled a Technological Sovereignty Package — a Chips Act 2.0 and a Cloud and AI Development Act — to make Europe “an AI continent” and “strengthen its digital autonomy.” And, as the closing pages record, a single US directive has just cut the most capable models off from every foreign national on Earth. The realist account is correct about all of this; it describes *what is*. This essay runs in a different lane — *what should be*, and what rightful arrangement could answer the order now emerging. The two are not opposites. They are complementary readings of the same hard fact, and the descriptive lane is what makes the normative one urgent.

Where this essay goes further is in two directions that frame leaves open. The first is downward — toward a kind of sovereignty that has nothing to do with dominance and does not require winning anything. The second is more uncomfortable, and I want to take it first, because it is the part Western readers are least prepared for.

The gaze we have not learned to turn

Audiences in the OECD democracies have been trained, carefully and for years, to mistrust the Chinese state’s relationship to technology. That training is not baseless. China’s National Intelligence Law obliges organisations and citizens to “support, assist and cooperate with” state intelligence work, and the implications for any data or model under Chinese jurisdiction are real and widely understood.

What we have *not* been trained to do is apply the same scrutiny, evenly, to the United States — and at this moment the symmetry deserves stating plainly, by someone who intends the United States no hostility whatever. Two American statutes reach directly into data held by American providers, wherever in the world it sits and whoever it belongs to. The **CLOUD Act** (2018) compels US-based technology companies to disclose data they control “regardless of whether

the data are stored in the U.S. or on foreign soil.” **Section 702** of the Foreign Intelligence Surveillance Act authorises warrantless surveillance of non-US persons located outside the United States, through the compelled assistance of US communication providers — a power reauthorised in April 2024 and, as this is written, up for renewal again in April 2026. These are not fringe readings. The European Data Protection Supervisor called the CLOUD Act “a law in possible conflict with the GDPR.” Germany’s data-protection commissioner warned government bodies against entrusting sensitive data to US cloud services. And — the comparison is not mine to soften — commentators noted that the CLOUD Act invited direct comparison to China’s National Intelligence Law.

The point is not that the United States is China, or that the two regimes are morally equivalent; they are not, and the differences in transparency, courts, and recourse matter. The point is narrower and sturdier: **for a community that is neither American nor Chinese, dependence on either sovereign’s stack is a standing exposure to a decision it has no part in.** The reflex to examine only one of the two superpowers is not even-handedness; it is a blind spot, and a blind spot is a place where authority quietly leaks away.

I will give this its hardest, most current form below, because it happened while this essay was being written, and because it happened to the very tool I am writing it with. First, though, the constructive turn — what kind of sovereignty answers the problem that neither patron can.

Sovereignty and legitimacy — two measures, not one word

There is a cost to the loose way the word has been used so far, and it is worth paying off. *Sovereignty*, used carefully, names one thing: an actor’s control over its own domain, to the exclusion of others. *Legitimacy* names a different thing: the rightful authority an actor holds in the eyes of the people it affects, within its borders and beyond them. They are separate axes, and collapsing them into a single contested word — as the “sovereign AI” debate tends to, and as an earlier draft of this essay did — throws away most of the information. (I owe the sharpening of this distinction to Tim Clancy of UMD ARLIS, who pressed it in correspondence after the first draft.)

Held apart, the two axes describe a continuum the race cannot see. An actor can hold near-total sovereignty and almost no legitimacy — full control of its domain that few, inside or out, regard as rightful. Another can hold little sovereignty and deep legitimacy — scant power to exclude, but wide recognition of its right to decide. The world’s hundred-and-eighty-odd state and emerging-state actors are scattered across that plane; the single-axis question — *who is winning the AI race* — flattens every one of them onto the same line.

And the race measures only that one axis, at its most punishing end. A serious recent assessment — the Brookings–CEPS study of February 2026 — concludes that “full-stack AI sovereignty is structurally infeasible for almost any country, because AI is a transnational stack with concentrated choke points,” and

warns that the national “sovereign AI” reflex “can also become a vehicle for protectionism, fragmented markets and standards, and duplicative or stranded public investment.” That is sovereignty-as-capacity at national scale — what NVIDIA’s Jensen Huang means by “every country needs to own the production of their own intelligence” — and most actors will never hold it.

Graylin’s remedy corrects the *direction* of that capacity race — cooperate, do not dominate — but it relocates authority no further than the summit; a “CERN for AI” is a convocation of the powers large enough to be invited. The two measures that his frame — and the race — leave unread are exactly the ones most of the world can move on: **legitimacy**, the rightful authority of the governed; and, beneath the nation, a modest but real **sovereignty** a community *can* hold — control over its own data and the models that act on it, to the exclusion of outside override. Neither requires winning the capacity race. Both can be held today by actors who will never own a single frontier cluster.

What the Village holds

The Village’s position is a coordinate on that map — high on legitimacy, modestly but genuinely sovereign over its own domain — and it can be stated in one sentence:

The sovereignty worth wanting is a community’s control over the data and systems that act on it, exercised with legitimacy — held at human scale, federated, never surrendered.

It rests on two things, held co-equally and exercised together as **co-governance**:

Legitimacy — the recognised, rightful standing of a community (a village, an iwi, a hauora, a municipality, a congregation) to make binding decisions about the AI that operates on it: what it may do, what it must refuse, who answers for it. Rightful authority in the eyes of the people it serves — not capacity.

Data sovereignty — control, at community scale, over the one domain a community can actually hold to the exclusion of outside override: the data it creates and the models that act on it. Data is not an asset extracted from a place; it is *of* the place, and stays under the place’s authority across its whole life. This is the sense in which Māori Data Sovereignty — *rangatiratanga* over data — already uses the word: sovereignty in the strict sense, scoped to the domain a community can govern rather than to territory it cannot.

Neither is subordinate to the other, and the conjunction matters. Legitimacy without data sovereignty is a seat at a table where someone else holds the records. Data sovereignty without legitimacy is control no one recognises as rightful. Together they describe a standing that is *not dominance over anyone* — self-determination plus the capacity to say no — and it scales not by conquering upward but by **federating** outward: many communities holding their own authority, reaching one another through consent-bound, revocable chan-

nels, rather than racing toward a centre. A federated mesh has no node that must win.

This is old, not invented

It would be a kind of arrogance to present this as a new idea. It is a very old one, recovered, and it appears across traditions that have nothing to do with one another — which is part of why it is worth trusting.

Elinor Ostrom won the 2009 Nobel Memorial Prize in Economics for demonstrating, against the prevailing fatalism of the “tragedy of the commons,” that communities can and do govern shared resources sustainably — a third way between privatisation and central control. Her field studies of irrigation systems, fisheries, and mountain commons yielded eight design principles for durable self-governance: clear boundaries, rules that fit local conditions, collective decision-making by those affected, monitoring, graduated sanctions, accessible conflict resolution, recognition by outside authorities of the community’s right to organise, and — for anything large — **nested enterprises**, governance layered from the local upward. Federation, in other words, is not a software pattern we invented; it is what enduring commons have always looked like.

Subsidiarity says the same thing in a different idiom. In its classic formulation it is “an injustice and at the same time a grave evil and disturbance of right order to assign to a greater and higher association what lesser and subordinate organizations can do” (*Quadragesimo Anno*, 1931); the principle is written into European Union law, where higher levels act only when they can do so more effectively than national or local ones. Authority belongs at the most local competent level. Everything above exists to serve it.

Simone Weil, writing in 1943, put it as a matter of the soul rather than of institutions: obligations come before rights, and rootedness — “real and active participation” in a living community that keeps the past alive and the future in view — is a basic human need, its loss a genuine harm. She was suspicious of rights talk precisely because it leans on force; she would ask, of any AI system, not what claims it might one day make but what is *owed* to the vulnerable realities already in its care.

And **te ao Māori** offers, in *rangatiratanga*, a deeply worked account of relational authority — the inherent right of a people to exercise control over what is theirs — given concrete expression in Māori data sovereignty, which holds that “Māori have an inherent right to exercise control over Māori data and Māori data ecosystems,” grounded in Te Tiriti o Waitangi. I name this tradition with deliberate care and as one among several, not as the property of this essay. Te Tiriti carries real weight in New Zealand law; the wider indigenous-rights instruments it is often paired with do not, and I will not overstate them. The reason to invoke it here is not to borrow its authority but to acknowledge that some of the clearest thinking about holding rightful authority over a domain

was done long before the present moment, by people the capacity race would never have counted.

Four traditions, no shared lineage, one conclusion: authority is most legitimate, and most durable, when it is held close to the ground and federated upward by consent.

The architecture: an immutable layer of constitutional values

A principle that lives only in prose is a sentiment. What makes “never surrendered” a property of a system rather than a promise is architecture, and this is where the Village’s claim becomes concrete and falsifiable.

Governance in the Village is layered, in a deliberate order. At the base sits an **immutable layer of constitutional values** — a small set of inviolable commitments, enforced *structurally* in the architecture itself rather than written into a policy document that can quietly drift. A few of the commitments that live at this level:

- certain decisions always require a human;
- when a community’s own cultural authority marks a piece of content as restricted, the model must refuse it or escalate it, rather than answer on its own;
- a person’s authority over their own data cannot be revoked without their knowing;
- Te Tiriti, where it applies, sits here — as a constitutional commitment, not a courtesy line added at the end.

Immutable means exactly that: no single party — not the platform operator, not a community, not the model — can override these commitments unilaterally, or erode them by stealth. A community may make the floor stricter for itself; none can be forced below it.

Above that floor, each community writes and edits its own rules: its kaupapa, its red lines, its permitted uses. Governance precedes infrastructure — the community sets the constraints, and the system is built to obey them, not the other way round. And the authorities involved — platform, iwi, community trust — are **co-equal peers**, each publishing its rules and able to withdraw them at any moment; when an authority withdraws, the system must stop relying on it. Legitimacy is composed from co-equal parts, never concentrated in one.

Co-equal governing principles, set in an immutable constitutional layer, are not decoration. They are the structural answer to the failure mode the next section describes — the slow normalisation by which good values, left to discretion, are eroded one reasonable-seeming step at a time. Values that cannot be quietly overridden cannot be quietly normalised away.

Why now: normalisation, and the closing window

Hannah Arendt’s most enduring and most misread observation is that great harm rarely arrives announcing itself as monstrous. It arrives as the ordinary and the administrative, accommodated by people exhibiting what she called “a curious, quite authentic inability to think” — not stupidity, not even conviction, but a failure to look at the world from any standpoint but one’s own. Evil of this kind, she wrote, “can overgrow and lay waste the whole world precisely because it spreads like a fungus on the surface.” (She did not, it must be said, excuse anyone; she held the individual fully responsible. The banality is in the mechanism, not the culpability.)

Turn this on the present. As an extractive, dominance-shaped mode of AI becomes ubiquitous and convenient, the public’s sense of what is acceptable is not overturned by any decision anyone defends. It is worn smooth by habituation. What should unsettle becomes background; guardrails erode not by repeal but by familiarity; and the whole system drifts, step by reasonable step, toward configurations less and less averse to catastrophe — with no one having chosen the destination. This is the danger the immutable constitutional layer is built against: it removes the most important values from the reach of drift.

And here the timing cuts both ways, which is the whole force of the argument. The same exponential rise in capability that accelerates the drift has, for the first time, made the alternative cheap to build. The shift is not subtle, and the clearest evidence is a demonstration rather than a recollection: the source-checking behind this essay was carried out by an automated research run that marshalled **more than a hundred specialised sub-agents in a single pass and finished, in minutes, work that a human research team would measure in days** — and it did so twice, to verify itself. In roughly a year, AI engineering has gone from an assistant that completes a single function to a system that orchestrates a swarm against a whole problem. The practical consequence is decisive: a small, principled team — even one whose own technical skills are modest — can now stand up a sovereign, multi-region system, running situated small models (Llama, Qwen) under its own governance on infrastructure outside the reach of any single state, in months rather than years. The thing that makes the danger faster is the same thing that finally arms the response. There is no longer an excuse to wait, and not much time in which waiting is safe.

A live exhibit, and an admission

I had intended to argue the dependency point in the abstract. Instead it made itself, in the days this essay was written — and at my own expense, which is the most telling way it could have been made.

On 12 June 2026, Anthropic announced that “the US government, citing national security authorities, has issued an export control directive to suspend all access to Fable 5 and Mythos 5 by any foreign national, whether inside or outside

the United States, including foreign national Anthropic employees.” Fable 5 was, at that point, the most capable model available — the very swarm-orchestrating capability described above. Overnight, by an order I had no standing in, it was withdrawn from every foreign national on Earth. The stated basis was a narrow, non-universal vulnerability — “asking the model to read a specific codebase and fix any software flaws” — a capability Anthropic itself observed is “widely available from other models (including OpenAI’s GPT-5.5).” Anthropic, to its considerable credit, disagreed with the directive in public and is working to restore access. The company is not the villain of this story. There is no villain. There is only the structure.

And the structure is exactly the one this essay describes. A capable practitioner in New Zealand, building tools for communities, had his most powerful instrument removed by a sovereign he does not answer to, for reasons of that sovereign’s national security, with no recourse available to him. That is not a hypothetical risk in a policy paper. It is sovereignty-as-dominance reaching down through the layers and touching a single person — and it is precisely the dependency the Village posture declines. It is also why the response has to be architectural rather than rhetorical. Three concrete forms answer a switch that can be flipped from elsewhere: **situated small-model inference** a community runs on its own infrastructure, so the capability does not live behind someone else’s directive; **federation** that reaches other communities through consent-bound, revocable channels rather than a single provider’s centre, so no one node can be cut off and the mesh fail; and **data-record sovereignty** — provenance and authority travelling with each record — so that even when a model is borrowed, the data and the right to govern it never leave the place. That is what “never surrendered” means once you accept that any single supplier can be ordered, overnight, to cut you off.

I should be plain about the reflexivity rather than hide it: this essay was drafted with the help of Anthropic’s models, on American infrastructure. The tool used to argue for a sovereign stack is itself subject to the dominance logic the stack is built to resist. That is not an embarrassment to be concealed; it is the argument, demonstrated. One governs the layer one can actually hold, and names the layer one cannot yet. Reducing that dependency — not pretending it away, not attacking those who supply it — is the whole of the practical programme.

The aspiration, stated plainly

I will not dress this up. The aspiration is that **every community should be able to run AI that it governs, and own the data it creates** — while honouring Te Tiriti and other co-equal governing principles, encoded in an immutable layer of constitutional values that no operator, no community, and no model can quietly override. Not every community winning the capacity race; almost none will. Every community holding the authority it is entitled to over the systems that act on it, at human scale, federated with others, and never surrendered.

That is an ambition, and I hold it firmly. It is not a boast. The Village does not claim to have solved AI governance, to be a rival to anyone’s frontier lab, or to have escaped the dependencies it openly names. It claims something smaller and harder to dismiss: that the layer where rightful authority can actually be held is buildable now, for actors who will never hold the other kind, and that building it is a quiet act of non-alignment — owing allegiance to no patron because it does not play the patrons’ game.

For different readers

This essay is written first for those who think about AI governance and policy, but the principle it defends is not theirs alone. A short translation for the others it concerns:

- **For philosophers:** the move is to dethrone consciousness and capacity as the master questions and restore *authority* and *obligation* to the centre — to ask, before “what can this system do,” the older questions of who may rightfully decide over it and what is owed to those in its care.
- **For systems architects and practitioners:** the same exponential capability that threatens to normalise an extractive AI is the capability that lets a small team build a governed, federated, situated-inference alternative on a sovereign stack — and the engineering discipline is to make the constitutional values *structural*, enforced in the architecture, not left to policy.
- **For communities and those who would carry this to them:** sovereignty here does not mean building your own data centre or matching anyone’s compute. It means holding rightful authority over the AI that acts on you, owning the data you create, and retaining the capacity to say no — and being able to do so alongside others without surrendering it to any of them.
- **For Māori governance bodies:** this is offered as a sketch from outside, naming *rangatiratanga* and Te Tiriti with care and as one tradition among several, never spoken for and never absorbed — and resting on an architecture in which Te Tiriti is a constitutional commitment, not a clause.

The posture throughout is meant to be the one the subject deserves: respectful of every geopolitical interest, hostile to none, and steadfast on the single principle worth being steadfast about. Before asking whose AI to build on, ask what authority you are entitled to hold yourself — and then hold it.

Sources

On the race and its critique

- Alvin W. Graylin & Paul Triolo, “There can be no winners in a US-China

AI arms race,” *MIT Technology Review* (Jan 2025). <https://www.technologyreview.com/2025/01/21/1110269/there-can-be-no-winners-in-a-us-china-ai-arms-race/>

- Alvin W. Graylin, “Beyond Rivalry: A US-China Policy Framework for the Age of Transformative AI,” *The Digitalist Papers*, Vol. 2 (Dec 2025). <https://www.digitalistpapers.com/vol2/graylin>
- Alvin W. Graylin, “America Is Running the Wrong AI Race” (Jul 2025). <https://abundanist.substack.com/p/america-is-running-the-wrong-ai-race>

On the race as “what is emerging” (the descriptive lane)

- Iranian drone strikes on Amazon (AWS) data centres, UAE & Bahrain, March 2026 — the first known military strikes on a US hyperscaler’s infrastructure. *Rest of World*. <https://restofworld.org/2026/iran-amazon-data-center-strikes/>
- European Technological Sovereignty Package (Chips Act 2.0; Cloud and AI Development Act) — “an AI continent,” “strengthen its digital autonomy.” European Commission, 3 June 2026. https://commission.europa.eu/news-and-media/news/strengthening-europes-tech-sovereignty-2026-06-03_en

On sovereignty as capacity, and its limits

- “NVIDIA CEO: Every Country Needs Sovereign AI,” NVIDIA blog (Feb 2024). <https://blogs.nvidia.com/blog/world-governments-summit/>
- “Is AI Sovereignty Possible? Balancing Autonomy and Interdependence,” Brookings / CEPS (Feb 2026). <https://www.brookings.edu/articles/is-ai-sovereignty-possible-balancing-autonomy-and-interdependence>

On state reach into data — the even-handed record

- CLOUD Act (2018) — extraterritorial reach over data held by US providers; EU/GDPR conflict; comparison to China’s National Intelligence Law. https://en.wikipedia.org/wiki/CLOUD_Act
- Section 702, Foreign Intelligence Surveillance Act — warrantless surveillance of non-US persons abroad; reauthorised Apr 2024, sunsets Apr 2026. https://en.wikipedia.org/wiki/Foreign_Intelligence_Surveillance_Act

The traditions

- Elinor Ostrom, *Governing the Commons* (1990); 2009 Nobel Memorial Prize in Economics; the eight design principles. https://en.wikipedia.org/wiki/Elinor_Ostrom
- Pope Pius XI, *Quadragesimo Anno* (1931), §79 — subsidiarity; and TEU Art. 5(3). https://www.vatican.va/content/pius-xi/en/encyclicals/documents/hf_p-xi_enc_19310515_quadragesimo-anno.html
- Simone Weil, *The Need for Roots* (1943/1949) — obligations before rights; rootedness. <https://plato.stanford.edu/entries/simone-weil/>

- Te Mana Raraunga, “Principles of Māori Data Sovereignty” (2018); the Charter. <https://www.temanararaunga.maori.nz/tutohinga>

On normalisation

- Hannah Arendt, *Eichmann in Jerusalem* (1963) and *The Life of the Mind* — the banality of evil as thoughtlessness; *Stanford Encyclopedia of Philosophy*, “Hannah Arendt.” <https://plato.stanford.edu/entries/arendt/>

The live exhibit

- “Statement on the US government directive to suspend access to Fable 5 and Mythos 5,” Anthropic (12 Jun 2026). <https://www.anthropic.com/news/fable-mythos-access>

The Village platform and the Tractatus framework are an attempt to make governance achievable for communities at human scale — relocating authority to where it can rightfully be held, and letting communities federate that authority rather than surrender it. This essay is offered in that spirit, and with respect to all whose interests it touches.

Copyright © 2026 John G. Stroh / My Digital Sovereignty Ltd. Licensed under CC BY 4.0 (Creative Commons): you are free to share and adapt this work, with attribution.