The Map Has No Node for Legitimacy

A response to Tim Clancy and Asmeret Bier Naugle's qualitative model of AI sovereignty — and why some of us are building for the scenarios it leaves out

Tim Clancy and Asmeret Bier Naugle's new model of AI sovereignty as an instrument of national power is one of the clearest pictures we have of the strategic-competition view of this field. It is worth taking seriously — which is why it is worth disagreeing with carefully. The model treats national power in agentic AI as a contest over five growth levers a country accumulates: accelerators, datasets, skilled workforce, electricity, and water. Each has its own ceiling, and each can be attacked — an adversary either backfills its own supply or degrades a rival's, whether by force, like Iran's March 2026 strikes on data centres in the UAE, or by cyber and economic coercion. As a description of how great powers are actually behaving, it is largely right. As a definition of what sovereignty *is*, it is incomplete — in a way that matters enormously for everyone who is not a great power.

What the model gets right

The realist core is sound, and the loop-by-loop construction is genuinely useful. Its best insight is that adversary action shows up as one thing — *degradation of availability* — across several levers at once: workforce, datasets, accelerators, and the physical substrate. That single frame unifies export controls, talent restrictions, and supply-chain attacks, which are usually treated as separate problems. Naming all five levers, rather than stopping at the familiar water-and-electricity argument, is the right level of detail. And the authors' supplementary — a beginner's reference on how agentic compute is measured and what drives data-centre demand — is a real contribution in its own right. My quarrel is with the model's *definition* of sovereignty, not its engineering.

Where it is incomplete — and the incompleteness is in the definition, not a missing variable

The paper defines AI sovereignty as the extent to which a nation independently controls its AI technologies. So every one of the five levers is a measure of *capacity* — something a state accumulates or loses. Nowhere is there a term

for legitimacy, for the authority of the people whose data trains and steers the system, or for whether the governed consent to how a model acts on their behalf. These are not missing numbers the model could add later. They are not in its vocabulary at all, because the definition is about control-as-capacity, not authority-as-right.

This is not a quibble about a missing box. It is a different theory of what sovereignty means. The model defines it as *capacity to independently control and out-compete*. There is an older and more durable definition: sovereignty as *rightful authority over a domain* — the recognised standing to make decisions others are bound to accept. A nation with vast compute but no legitimate authority over how that compute is used is powerful, but not obviously sovereign in this second sense. A community with modest compute but genuine, recognised authority over its own data — and over how models steer on its behalf — may be sovereign in a way no lever count can measure.

Here is a useful test. Ask the model where the European Union's domestic AI-sovereignty measures belong on its map. Or where Māori data sovereignty belongs — *rangatiratanga*, authority and control over data, in Te Mana Raraunga's sense. Or any indigenous governance claim. There is no lever for any of them. They are neither adversary degradation nor capacity accumulation, so the model cannot represent them at all. The EU's June 2026 technological-sovereignty package is exactly this kind of instrument — a live assertion of AI sovereignty that a capacity model has no way to score. And these are precisely the tools by which smaller actors are asserting sovereignty right now: not by out-building Nvidia clusters, but by establishing rightful authority over data, provenance, and steering.

Two concepts of sovereignty

The disagreement is old — it predates AI by centuries. Political and legal theory has long separated two things the word "sovereignty" runs together: *effective control*, the de facto power to act without being overridden, and *rightful authority*, the de jure standing to make decisions others are bound to recognise. A state can hold one without the other. A junta has control without authority; a government-in-exile has authority without control. The two usually travel together, which is why the distinction is easy to forget. But they are not the same thing, and AI is prising them apart.

Clancy and Naugle measure the first. Their five levers are quantities of capacity — the de facto ability to build, run, and defend AI without depending on anyone. That is a real thing to measure. But the model is silent on the second. Nothing in it speaks to whether a system's behaviour is *authorised* by the people it acts upon, or whether the governed have any standing to bind it. In the older legal tradition, that standing is exactly what sovereignty *was*: not the largest army, but the rightful title to decide within a domain.

The modern default is the Westphalian settlement — sovereignty as territorial,

capacity-backed, and non-interfering. It is a poor fit for data, and a poorer one for peoples whose authority was never territorial in the first place. Māori data sovereignty is not a bid to out-compute anyone. It is *rangatiratanga* — authority grounded in relationship and whakapapa, a claim that holds whether or not capacity backs it. The European Union's measures are, likewise, assertions of rightful authority over how AI is governed inside a jurisdiction, not entries in a capacity race. Read as capacity claims they look weak; read as authority claims they are exactly what they are. The governance form that fits this picture is not a single sovereign but a polycentric one — many co-equal authorities, each with its own jurisdiction over a shared resource, in Ostrom's sense. That, not by coincidence, is the architecture the last section describes.

This is why the capacity definition is not neutral. It does not merely leave authority out — it makes authority impossible to state, because everything the model can express is counted in units of capacity. Once you separate the two concepts, the rest of this argument follows: a community can hold rightful authority over its own data, and over how a model acts on its behalf, without having the capacity to build that model. Authority is the layer that does not require winning the race.

The trap the model describes

Here is the part the model shows better than its authors may intend. On its own axes, the contest is unwinnable for anyone without national-scale capital. A small nation will never accumulate competitive zettaFLOPS, never out-build accelerator supply, never win a water-and-electricity war of attrition against a continental power. The authors say plainly that each lever's limits are set by a country's own capabilities and by its ability to call on partner nations to backfill what it lacks — and that is precisely the structural trap. If sovereignty *is* capacity across these five levers, then most of the world is permanently dependent, and the only rational move is to pick a patron and accept it.

That conclusion should make us suspect the premise. A definition of sovereignty that permanently locks out most nations and communities is not a neutral measurement. It is also, as Yew and colleagues argue in their work on the commodification of AI sovereignty, the move that lets private vendors define the term on their own terms: once "sovereign" AI factories, clouds, and models are things you can buy, the firms that sell them get to say what counts. Measuring sovereignty in units of capacity is what makes that capture possible — because those units are exactly what they sell. When sovereignty is counted in accelerators and zettaFLOPS, the companies that sell accelerators and zettaFLOPS get to decide who is sovereign.

Where the Village fits — and why it can thrive in the gaps

My own work — the Village multi-tenant platform and the Tractatus governance framework, built under the My Digital Sovereignty initiative — starts from the

second definition, the one the model has no lever for. The claim is concrete, and it is running in production. Every record carries its own provenance: a cryptographic origin hash and a tamper-evident, append-only log of who did what to it. A community body — an iwi, a marae, a club, a whānau — holds real, co-equal authority over how the model behaves on its data. It is served by a model fine-tuned to its kaupapa. It writes and edits its own governance rules, layered over a platform safety floor it can raise but can never be forced below. Through its own cultural authority it can mark content as tapu or restricted, and the model must then refuse it or escalate to that authority rather than answer on its own. The authorities themselves — platform, iwi, community trust — are treated as co-equal peers. Each publishes the rules and steering that shape the model's behaviour in its own jurisdiction, and each can withdraw them at any moment, at which point the platform must stop using them. None of this is a roadmap item. It is deployed, composed at inference, and auditable end to end. It is sovereignty as rightful authority, asserted at a layer that does not require winning the lever race.

I want to be precise about what is and is not solved here, because the incompleteness cuts both ways. The Village governs the layer it can actually govern. The substrate beneath it — the frontier model weights, the accelerators, the compute itself — sits inside exactly the foreign-controlled levers Clancy and Naugle map. The models we run are open weights authored outside our jurisdiction, served locally on EU- and NZ-hosted GPUs. I do not pretend that dependency away. My claim is narrower: governance sovereignty and substrate sovereignty are separable, and the first is built and running today for actors who will never hold the second. The model itself shows *why* substrate sovereignty is out of reach for small actors — it sits downstream of national-scale resource loops they do not control. This fits Singh and Sengupta's proposal to treat sovereignty as a continuum rather than a binary: autonomy at one layer, acknowledged dependence at another. That honest separation is the whole architecture.

The "third option" follows directly. The strategic-competition frame offers small nations a binary: align with American big-tech AI or Chinese big-tech AI, and accept the dependency either way. But if sovereignty is rightful authority rather than raw capacity, a third path opens that depends on neither — not by matching their levers, which is hopeless, but by refusing their definition of the contest. The actors the model's loops exclude are not powerless. They are being measured on the wrong axis.

The friendly amendment

So this is less a rebuttal than a request to widen the model. Add the nodes the realist frame leaves out: legitimacy, data-subject authority, the standing of the governed to steer. Once those are in, a second set of loops appears — ones where a community can gain sovereignty without gaining a single accelerator, and where an adversary's compute advantage does not automatically erode another actor's rightful authority over its own domain. Those loops are where much of

the world will have to live, because the capacity loops were never open to it in the first place.

The strategic-competition view is a true account of how the powerful are behaving. It is not a complete account of what sovereignty is. And in the gap between those two things — the scenarios the model cannot represent — is exactly where the rest of us are building.

References

- Clancy, T. & Naugle, A. B. (2026). *AI Sovereignty: A Qualitative Model of Strategic Competition as AI Becomes an Instrument of National Power.* arXiv:2606.07245.
- Yew, R.-J., Creasey, K. E., Curtis, T. L. & Venkatasubramanian, S. (2026). *The Commodification of AI Sovereignty: Lessons from the Fight for Sovereign Oil.* arXiv:2601.11763.
- Singh, S. K. & Sengupta, S. (2025). *Sovereign AI: Rethinking Autonomy in the Age of Global Interdependence.* arXiv:2511.15734.
- Te Mana Raraunga — Māori Data Sovereignty Network (2018). *Principles of Māori Data Sovereignty.*
- Kukutai, T., Campbell-Kamariera, K., Mead, A., Mikaere, K., Moses, C., Whitehead, J. & Cormack, D. (2023). *Māori Data Governance Model.* Te Kāhui Raraunga.
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action.* Cambridge University Press.

The Village platform and the Tractatus framework are an attempt to make AI sovereignty achievable for actors who will never win the capacity race, by relocating sovereignty to the layer where rightful authority can actually be held.