



Computational Sovereignty and State Capacity
in the Age of Artificial Intelligence:
Toward a Quantitative Model for Measuring
Technological Dependence Across States

Dr. Hafid MRHIZOU

PhD in Legal and Political Sciences

Sidi Mohamed Ben Abdellah University – Fez

Morocco

Abstract

This article examines computational sovereignty and state capacity in the age of artificial intelligence by distinguishing the computational resources a state can use from the institutional control it retains over their use, continuity, and substitution. It argues that greater AI use does not automatically increase state capacity: its effects depend on data quality, AI governance and procurement capacity, internal technical expertise, and the structure of dependence on external vendors and infrastructure.

It introduces the Computational Sovereignty Gap as a diagnostic difference between use and control after both are calibrated to a common reference scale, and develops a quantitative framework linking usable computational capacity, sovereign computational capacity, AI governance capacity, data quality, and external technological dependence. Computational scenarios test the model's internal consistency and sensitivity, with particular attention to the Global South. The central claim is that computational sovereignty is not self-sufficiency, but the capacity to govern dependence while preserving meaningful options for audit, portability, modification, substitution, continuity, and exit.

Keywords: computational sovereignty; state capacity; artificial intelligence; technological dependence; Computational Sovereignty Gap; digital governance; Global South.



List of Abbreviations and Mathematical Symbols

List of Abbreviations

Abbreviation	Meaning
<i>AI</i>	Artificial Intelligence
<i>AIC</i>	Usable Computational Capacity
<i>AIGC</i>	AI Governance Capacity
<i>AU</i>	Understanding and Auditing
<i>CO</i>	Continuity and Exit
<i>CSC</i>	Sovereign Computational Capacity
<i>CSG</i>	Computational Sovereignty Gap
<i>DQ</i>	Data Quality
<i>ED</i>	External Technological Dependence
<i>HC</i>	Internal Human and Technical Expertise
<i>MS</i>	Modifiability and Substitutability
<i>PO</i>	Portability of Data and Functions and Interoperability
<i>SC</i>	State Capacity
<i>JRC</i>	Joint Research Centre
<i>OECD</i>	Organisation for Economic Co-operation and Development

List of Mathematical Symbols

Symbol	Meaning
<i>i</i>	Country index
<i>j</i>	Indicator index
<i>K</i>	Dimension/weight index, $k=1,...,5$
X_{ij}	Raw value of country <i>i</i> on indicator <i>j</i>
L_j	Lower reference bound for indicator <i>j</i>
U_j	Upper reference bound for indicator <i>j</i>
Z_{ij}	Normalized value of country <i>i</i> on indicator <i>j</i>
w_k	Weight assigned to dimension <i>k</i> of the <i>CSC</i> index
$\sum_{k=1}^5 w_k = 1$	Constraint requiring the <i>CSC</i> weights to sum to one
β_0	Intercept or baseline term in the regression model
$\beta_1,...,\beta_5$	Coefficients associated with the explanatory variables in the baseline model
β_6	Coefficient of the interaction term <i>AIC_i×Ed_i</i>
ε_i	Error term capturing factors not included in the model
<i>AIC_i×Ed_i</i>	Interaction between usable computational capacity and external technological dependence
$\partial SC_i / \partial AIC_i$	Marginal effect of usable computational capacity on state capacity
<i>ED*</i>	Threshold level of external technological dependence at which the marginal effect of <i>AIC</i> becomes zero in the linear interaction model
Σ	Summation operator
$\approx$	Approximately equal to
[0,1]	Common normalized range for the core model variables
[-1,1]	Theoretical range of the Computational Sovereignty Gap when AICAIC and CSC are bounded in [0,1]



Introduction

Artificial intelligence is becoming part of the infrastructure through which states process data, deliver services, support decisions, allocate resources, and manage risk. OECD evidence shows widespread adoption across government functions alongside substantial differences in data readiness, infrastructure, skills, and organizational capacity (OECD, 2025, p. 108). Yet wider use does not necessarily expand state control over the technical conditions that make such use possible. Government AI depends on interconnected resources—data, compute, cloud services, foundation models, software, expertise, and contracts—some of which public institutions may access without owning or fully controlling.

This distinction reaches beyond administrative efficiency to the distribution of power in the digital economy. Farrell and Newman show that actors controlling central network nodes can impose costs on more dependent parties (Farrell & Newman, 2019, p. 45). In government computing, dependence therefore turns not only on where a resource originates, but on whether realistic alternatives exist. Roberts similarly argues that digital sovereignty concerns not only *de facto* control but legitimate authority over technology (Roberts, 2024, p. 4), a distinction that matters when private infrastructures support continuous public functions for which the state remains politically accountable.

The article therefore narrows the broad idea of digital sovereignty to computational sovereignty: effective institutional control over the resources required to operate AI, rather than ownership of any single component. A state may control its data legally while relying on cloud infrastructure, models, or expertise it cannot independently modify; conversely, it may rely on external services while retaining audit, portability, substitution, and exit rights. External provision is not inherently a weakness and may lower costs or accelerate access to advanced capabilities. The problem arises when difficult-to-substitute vendors, closed architectures, or weak procurement capacity constrain institutional choice. OECD reviews associate weak AI-procurement preparedness with risks including vendor lock-in, unclear accountability, limited transparency, and data-governance problems (OECD, 2025, pp. 108, 116).

This produces a central paradox: AI can raise operational capacity while reducing sovereign control over the infrastructure on which that capacity depends. A state may perform some functions more effectively yet become more exposed to vendors, platforms, or expertise it cannot readily replace. This problem bears directly on state capacity, understood here in Fukuyama's sense as the strength and competence of the administrative apparatus to implement policy rather than the breadth of state functions alone (Fukuyama, 2013, p. 350).

The article distinguishes usable computational capacity—the computational and algorithmic resources a state can access and deploy—from sovereign computational capacity—the institutional control it retains through understanding and auditing, portability and interoperability, modification and substitution, internal expertise, continuity, and exit. Computational sovereignty therefore does not require technological self-sufficiency. It denotes the ability to use computational resources while preserving effective authority over the conditions of their use, including the ability to review, change, or terminate technological arrangements when public objectives require it (Roberts, 2024, p. 4).

The core hypothesis is that AI's effect on state capacity is conditional rather than linear. Advanced systems are more likely to produce durable gains when paired with good data, governance capacity, technical expertise, and viable alternatives; where external exposure is concentrated and review or substitution is weak, operational gains may coexist with institutional vulnerability. This issue is especially salient in the Global South, where external infrastructure may overcome constraints of capital, expertise, and time without implying that these countries form a homogeneous or necessarily dependent group.

The central research question is: how does the gap between the computational capacity a state uses and the capacity it can control affect state capacity, and when does reliance on external AI infrastructure shift from a source of increased capacity to a source of institutional dependence? The



article contributes by distinguishing usable from sovereign computational capacity, proposing the Computational Sovereignty Gap as a diagnostic standardized difference between use and control, and developing a quantitative model of use, control, dependence, and state capacity. The gap is treated descriptively or in an alternative specification, not as an additional regressor alongside the components from which it is derived.

The article proceeds in six sections. Section One develops the concept of computational sovereignty; Section Two disaggregates the infrastructure of computational capacity; Section Three analyzes technological dependence; Section Four formalizes the quantitative model; Section Five tests its internal consistency through scenarios and simulation; and Section Six considers its implications for the Global South.

Section One: From Digital Sovereignty to Computational Sovereignty

This section defines computational sovereignty by separating formal ownership from effective institutional control. It then identifies the dimensions of control that make external dependence governable and links them to state capacity and the Computational Sovereignty Gap.

1. From Digital Sovereignty to the Concept of Computational Sovereignty

Digital sovereignty increasingly concerns control over the infrastructures, technologies, and data through which states exercise public authority. Couture and Toupin show that the term encompasses multiple claims to independence, control, and autonomy across states, firms, users, and infrastructure operators (Couture & Toupin, 2019, p. 2306). The issue becomes more specific when digital infrastructure forms part of executive capacity: legal jurisdiction over a service does not necessarily entail technical control over the systems through which it is delivered.

Roberts distinguishes *de facto* control from legitimate authority over technology, emphasizing that sovereignty also concerns who may justifiably determine how technology is used and with what consequences (Roberts, 2024, p. 4). This is particularly relevant to AI, where private firms may control key components while the state remains accountable for dependent services. Comparative evidence likewise shows that declared sovereignty policies may diverge from actual hosting and traffic patterns and involve trade-offs among autonomy, security, and efficiency (Jansen et al., 2023, p. 2). Computational sovereignty is therefore treated here as an observable property of infrastructure and practice, not a policy slogan.

2. Institutional Control between Ownership and Technical Authority

Ownership alone is an incomplete measure of control. A state may legally own data or a system yet depend on external cloud services, models, or expertise it cannot inspect or modify; conversely, it may use infrastructure it does not own while retaining contractual rights to audit, transfer data, change providers, and continue the public function. AI intensifies this problem because its operational chain spans data, compute, data centers, cloud services, models, software, interfaces, and human expertise, often under different actors' control (Roberts, 2024, p. 7). Data sovereignty therefore cannot capture dependence located in computing, software, or knowledge layers.

Computational sovereignty is consequently a specific dimension of digital sovereignty: the state's capacity to control the resources that make computational processing and AI systems possible and sustainable. The distinction matters as governments rely more heavily on cloud and shared services, since the ability to scale AI varies with infrastructure, skills, and organizational capacity and may depend on resources outside public institutions (OECD, 2025, p. 135).

3. Governing Dependence, Substitutability, and Continuity of Public Functions

External dependence is not a problem in itself. Interdependence is normal in the digital economy, and external services may be more efficient than underused public infrastructure. The political issue begins when substitution becomes difficult or exit costs materially reduce institutional choice. A relationship remains governable when data, applications, and operational knowledge can move



between providers; it approaches vendor lock-in when closed architectures, cumulative switching costs, or loss of internal expertise make exit impractical. OECD identifies such lock-in as a risk where government AI-procurement capacity is weak (OECD, 2025, p. 116).

Computational sovereignty therefore means governing dependence rather than eliminating it. A state may rely extensively on external vendors and still retain substantial sovereignty if it can understand, audit, alter, and terminate those relationships; local hosting alone does not guarantee equivalent control.

4. Dimensions of Computational Sovereignty and Institutional Measurement

Five dimensions translate this idea into institutional measurement: understanding and auditing; portability of data and functions and interoperability; modifiability and substitutability; internal human and technical expertise; and continuity and exit. These dimensions are graded rather than binary, because a state may control some layers strongly and others weakly. Computational sovereignty is therefore treated as a multidimensional variable rather than a distinction between 'sovereign' and 'non-sovereign' states.

The same dimensions later form the basis of the sovereign computational capacity index. Vendor concentration and external exposure are measured separately so that indicators of control are not duplicated across multiple variables.

5. Computational Sovereignty and State Capacity: From Dependence to the Sovereignty Gap

The framework also distinguishes reliance from dependence. Reliance means using resources supplied by another actor; dependence arises when alternatives, bargaining power, or exit options are sufficiently limited to constrain institutional choice. External infrastructure may therefore increase state capacity when viable alternatives remain available, while core functions built around a difficult-to-replace provider can create vulnerability even when operational performance is high.

This distinction underpins the later separation between usable computational capacity, which measures what the state can deploy, and sovereign computational capacity, which measures control over the conditions of deployment. Their standardized difference forms the Computational Sovereignty Gap.

Section Two: The Political Structure of Computational Capacity

This section disaggregates the infrastructure over which computational sovereignty is exercised into physical computing, cloud infrastructure, data, models and software, and human and institutional capabilities. The aim is to locate dependence and control within each layer rather than infer sovereignty from resource scale alone.

1. Physical Computing, Semiconductors, and Supply Chains

AI depends on a layered infrastructure of processors, data centers, cloud services, data, software, foundation models, and human skills. Control over applications therefore does not imply control over the underlying resources that make them possible (Roberts, 2024, p. 5). Computing should accordingly be treated as a distinct foundation of AI readiness alongside connectivity, context, and human capability, with attention to availability, access, and resilience rather than physical ownership alone (World Bank, 2025, p. 39).

The political value of computing is not determined by scale alone. A state may host substantial capacity while relying on imported processors, software, or updates supplied by a narrow set of firms; another may own little domestic infrastructure yet access several interchangeable providers. Analysis should therefore distinguish resource availability, access, and control over conditions of use. For semiconductors and specialized accelerators, the relevant issue is not whether every component is domestically produced but whether access channels are diversified and substitutes can be activated at acceptable cost and speed (World Bank, 2025, pp. 39, 47).



This emphasis follows the logic of asymmetric interdependence. International specialization becomes politically consequential when alternatives are concentrated around nodes that are difficult to bypass (Farrell & Newman, 2019, p. 45). Computational risk should therefore be assessed through substitutability, diversification, and transition costs rather than the share of imported components alone.

2. Cloud Infrastructure, Data Centers, and the Distribution of Control

Cloud computing separates use from ownership by making storage, processing, and software available on demand, lowering entry costs and facilitating scale. Yet responsibility and control vary by service model: moving from infrastructure-as-a-service toward software-as-a-service reduces the government's technical management burden while potentially increasing vendor lock-in because more components are controlled by the provider (World Bank, 2023, p. 17). Managed services can therefore improve efficiency and security, but only if public institutions retain enough expertise to evaluate providers and move to alternatives when necessary.

3. Data, Portability, and Interoperability

Data are a third layer of computational capacity. AI value depends not only on compute but also on the quality, diversity, organization, and local relevance of data; the World Bank frames AI foundations as connectivity, compute, context, and capability (World Bank, 2025). Legal ownership, however, does not guarantee control if data formats, interfaces, or integration processes make transfer costly. Portability and compatibility are therefore components of computational control, and cloud-procurement guidance recommends assessing migration costs before contracting because lower transfer costs reduce vendor lock-in (World Bank, 2023, p. 24).

Interoperability extends this logic. Open standards preserve the ability to reconfigure infrastructure without rebuilding services, whereas closed architectures raise switching costs as applications, data, and expertise become tied to a provider. Technical standards therefore shape the state's future option set.

4. Models, Software, and Vendor Lock-in

Models and software create a similar distinction between use and control. A public body may integrate an external model through an API and gain rapid access to innovation while becoming exposed to provider decisions on updates, pricing, or terms of service. Dependence becomes structural when data, software, integration, and expertise are tied to the same provider strongly enough that migration would disrupt the public function. It is therefore the depth of integration, not vendor expenditure alone, that determines the severity of dependence.

Vendor lock-in makes this condition observable. In cloud procurement, portability concerns the ease and cost of moving systems and data between providers; greater portability lowers lock-in risk, which is why migration costs should be incorporated into procurement from the outset (World Bank, 2023, p. 24). Exit is thus a measurable institutional capability, not merely a normative aspiration.

5. Human Capabilities, Institutional Coordination, and Global Inequality

Human and institutional capabilities are the final layer. A state may procure advanced systems yet lack staff able to specify requirements, assess risks, monitor performance, or negotiate effectively. Internal expertise need not replace external provision; it enables the state to govern it. Public institutions do not need to internalize every technical function, but they do need enough knowledge to decide what should be owned, procured, diversified, or retained as a strategic capability.

This creates a coordination problem. Multi-vendor, hybrid-cloud, or multi-cloud arrangements can reduce concentration but increase technical and organizational complexity. Sovereignty therefore cannot be equated with maximal technological independence: poorly maintained domestic systems may be less resilient than well-governed external services. The relevant criterion is whether institutional choices preserve continuity and future options.



Global inequalities make this issue comparatively important. AI-related data centers, cloud capacity, and computing resources are distributed unevenly, while measurement of the 'compute divide' remains difficult because comparative data are incomplete (World Bank, 2025). Unequal location alone, however, does not establish political dependence: a state with little domestic compute may remain resilient if several substitutes exist, while extensive local infrastructure may still depend on a single foreign supplier of a critical component.

For this article, computational infrastructure therefore comprises five interconnected elements: physical computing, cloud infrastructure, data, models and software, and human and institutional capacity. Weakness in one layer can constrain the value of the others, and ownership of several layers does not eliminate dependence if an uncontrolled component is an unavoidable bottleneck. The core proposition is simple: not every reliance is dependence, and not every form of ownership constitutes sovereignty.

Section Three: Technological Dependence as a Variable in State Capacity

This section examines when technological reliance becomes dependence that constrains state choice. It focuses on procurement, information asymmetry, exit costs, path dependence, supplier concentration, and bargaining power, and then specifies the conditions under which reliance ceases to be merely a source of capacity.

1. From Technological Reliance to Technological Dependence

Technological reliance exists whenever the state uses resources or services it does not produce internally; it is common and not inherently harmful. Dependence emerges when continuity, modification, or operating conditions become tied to an external actor in ways that materially reduce the state's ability to substitute, renegotiate, or exit. This distinction avoids a false opposition between internal development and external procurement: in-house systems may be unsustainable without skills and maintenance resources, while well-designed contracts can provide advanced capabilities without sacrificing audit, portability, or substitution rights.

2. Public Procurement, Information Asymmetry, and Mediated Capacity

AI procurement can deepen dependence because vendors often remain involved after contracting through maintenance, updating, retraining, and infrastructure changes. Hickok emphasizes that procurement therefore extends across the system life cycle and requires continuing technical oversight (Hickok, 2024, pp. 1218–1219). Information asymmetry intensifies the problem when public institutions lack the expertise to specify requirements or independently evaluate performance, allowing the supplier to influence the standards by which its own system is judged.

Busuioc similarly notes that public institutions may depend on technology firms for knowledge needed to understand and challenge algorithmic systems (Busuioc, 2021, p. 830). The capacity to contract is therefore not equivalent to the capacity to govern. External procurement may rapidly expand operational capacity, but this gain remains sustainable only if the state can define objectives, verify performance, revise contracts, and move to alternatives without ceding those functions to the vendor.

This distinction can be expressed as direct versus mediated capacity. Direct capacity consists of resources the state can mobilize internally; mediated capacity is accessed through markets, contracts, or partnerships. Neither is inherently superior, but mediated capacity becomes more fragile as alternatives narrow and exit costs rise.

3. Exit Costs, Path Dependence, and Portability

Exit capacity is central to distinguishing reliance from dependence. Legal termination rights matter little if the state cannot transfer data, restart applications, obtain documentation, or access an alternative provider. The World Bank therefore treats portability and migration costs as key mechanisms for reducing cloud lock-in and recommends assessing them from the procurement stage (World Bank,



2023, p. 24). Dependence can thus be understood partly as rising financial, organizational, and knowledge costs of changing course, even without direct coercion by a vendor.

These costs accumulate over time. As databases, interfaces, staff routines, and services become adapted to one system, early technical choices can narrow later alternatives. This path dependence means that current contracts alone are an incomplete measure of vulnerability; governance must preserve future options through open standards, interoperability, documentation, and architectures that keep transition feasible.

Dependence may also be primarily epistemic. A state can possess source code or data yet lack personnel able to understand or audit the system, making a theoretically replaceable vendor difficult to substitute in practice. Internal technical capacity therefore means retaining enough expertise to negotiate, evaluate, and supervise external provision.

4. Supplier Concentration, Bargaining Power, and Imported State Capacity

Technological dependence may link public institutions directly to transnational firms rather than only to other states. Major technology companies control important cloud, model, and software layers (Roberts, 2024, p. 7), and control over central network nodes can create asymmetric power without implying that coercion is always exercised (Farrell & Newman, 2019, p. 45). The relevant concern is exposure to decisions made outside the public institution when continuity depends on infrastructure it cannot readily bypass.

Market concentration alone is therefore not proof of lost sovereignty. Its significance depends on viable alternatives, interoperability, data rights, transition capacity, and the bargaining power of the public purchaser. States differ in procurement scale and legal and technical expertise, so the same product can generate different degrees of dependence depending on how the contractual relationship is structured and monitored.

This variation is especially relevant in the Global South, where unequal access to advanced compute and expertise may increase reliance on external resources without making dependence inevitable. Evidence from a Brazilian digital project likewise shows that organizational autonomy can persist within globally provided infrastructures when actors retain knowledge, alternatives, and control over objectives (Estima, 2026, p. 4). Structural dependence therefore constrains agency but does not mechanically determine it.

The idea of 'imported state capacity' captures situations in which a state expands functions through technical or knowledge resources produced outside the public apparatus. The capacity is real, but part of its sustainability depends on conditions the state may not fully control. This makes it useful to distinguish the volume of capacity, its source, and the degree of control over it—dimensions that can diverge even when observed administrative performance is similar.

5. The Dynamics of Dependence and the Conditions under Which Reliance Becomes a Constraint on Capacity

Dependence also has a temporal direction. External resources can reduce dependence when they transfer knowledge and help build alternatives, or deepen it when reliance persists without such learning. Localization is therefore not an automatic remedy: a locally hosted service may remain tied to external software, expertise, updates, and contracts, while weak domestic substitutes may reduce performance without increasing effective control. The relevant criterion is exercisable institutional control, not hardware location.

For this article, technological dependence is a condition in which the continuity or modification of public capacity becomes tied to external resources in ways that narrow effective alternatives. The quantitative model does not collapse this concept into one variable. It separates external exposure and concentration from the state's capacity to manage that exposure through auditing, portability, modification, substitution, internal expertise, continuity, and exit. Governed external reliance may



therefore expand capacity; dependence emerges when exposure combines with weak alternatives or weak sovereign capacity.

Operationally, *ED* captures concentrated or critical external exposure, whereas *CSC* captures the institutional instruments of control. *ED* is therefore one dimension of dependence rather than the whole concept; choice-constraining dependence becomes most consequential when high exposure coincides with weak sovereign capacity.

Section Four: Building a Quantitative Model of State Capacity and Computational Sovereignty

This section converts the theoretical framework into a testable quantitative structure. It separates use, institutional control, external exposure, and state capacity; defines the main indicators; specifies additive and interaction models; and sets out requirements for empirical identification and robustness.

1. Measurement Framework and Dependent Variable: State Capacity

Quantification requires conceptual discipline. The model therefore begins from the theoretical framework rather than from whatever indicators happen to be available, consistent with OECD/JRC guidance that conceptual definition should precede indicator selection (OECD & JRC, 2008, p. 22). Three analytical levels remain distinct: the technical capacity the state can use, institutional control over that capacity, and the political-administrative outcome represented by state capacity. This separation prevents AI diffusion from being treated as direct evidence of either state capacity or computational sovereignty.

This approach is consistent with Fukuyama's emphasis on administrative capacity rather than outputs alone, since observed outcomes may reflect social and economic conditions beyond government control (Fukuyama, 2013, p. 350).

State capacity in country i is denoted by:

SC_i

SC_i denotes the ability of state institutions to translate decisions and policies into effective implementation, not the size of the state or breadth of its functions (Fukuyama, 2013, p. 350). To avoid circularity, empirical measures of SC should exclude AI adoption itself; the outcome must be measured independently from the variables used to explain it.

2. Measuring Usable Computational Capacity and Sovereign Computational Capacity

Usable computational capacity is denoted by:

AIC_i

AIC_i measures computational and algorithmic capacity actually available for government use, whether owned, rented, or accessed through contracts. It may combine access to compute, cloud use, and diffusion of AI across government functions. Component weights should be justified theoretically and tested empirically rather than fixed arbitrarily; weighting and aggregation are themselves stages of composite-indicator construction (OECD & JRC, 2008, p. 89).

Sovereign computational capacity, the central control variable, is denoted by:

CSC_i

CSC_i measures effective institutional control over computational resources, not ownership alone. A public body may use external infrastructure while retaining strong control through audit, portability, substitution, and continuity rights, or own infrastructure while remaining dependent on external expertise and updates. *CSC* therefore comprises five dimensions: understanding and auditing; portability and interoperability; modifiability and substitutability; internal human and technical expertise; and continuity and exit (Roberts, 2024, p. 4).



Indicators for these dimensions must be placed on a common scale before aggregation. For cross-national and longitudinal comparison, fixed or theoretically defined reference bounds are preferable to sample-specific minima and maxima. Normalization is written as:

$$Z_{ij} = (X_{ij} - L_j) / (U_j - L_j)$$

X_{ij} is country i 's value on indicator j ; L_j and U_j are fixed lower and upper reference bounds; and Z_{ij} is the normalized value. Where theoretical bounds are unavailable, fixed empirical bounds from the full dataset or all available years can be used and varied in sensitivity tests.

The sovereign computational capacity index is then:

$$CSC_i = w_1AU_i + w_2PO_i + w_3MS_i + w_4HC_i + w_5CO_i$$

AU denotes understanding and auditing; PO portability and interoperability; MS modifiability and substitutability; HC internal human and technical expertise; and CO continuity and exit. Indicator direction is harmonized so that higher values consistently represent greater control.

Weights satisfy:

$$w_k \geq 0, \sum_{k=1}^5 w_k = 1$$

Equal weights provide a transparent baseline, followed by alternatives based on expert judgment or the data structure. Because different weights can change index values and country rankings, weighting and aggregation should be tested for sensitivity (OECD & JRC, 2008, p. 89).

3. The Computational Sovereignty Gap and External Technological Dependence

The Computational Sovereignty Gap is the standardized difference between use and control:

$$CSG_i = AIC_i - CSC_i$$

CSG_i is interpretable only when AIC_i and CSC_i are calibrated to a common reference scale with substantively consistent bounds. It is therefore a signed diagnostic indicator, not a physical measure of resources.

When the two values are close:

$$CSG_i \approx 0$$

institutional control is close to the level of use. If:

$$CSG_i > 0$$

use exceeds control, and the gap widens as that difference grows. If:

$$CSG_i < 0$$

control exceeds current use on the common scale; this is a relative control margin, not a physical reserve. With AIC and CSC bounded in $[0,1]$, CSG lies in $[-1,1]$.

For example, if $AIC = 0.85$ and $CSC = 0.45$:

$$CSG = 0.85 - 0.45 = 0.40$$

whereas if $AIC = 0.60$ and $CSC = 0.75$:

$$CSG = 0.60 - 0.75 = -0.15$$

The two cases differ politically even if the first state uses more advanced systems, because their levels of institutional control relative to use differ.

The model also includes an external-exposure variable:

$$ED_i$$

ED captures concentrated or critical external exposure, such as supplier concentration, the share of critical workloads dependent on external infrastructure or models, and concentration of access



channels. To avoid overlap with *CSC*, portability, substitutability, internal expertise, continuity, and exit are excluded from *ED*. High transition costs remain evidence of dependence and can be used for validation or case analysis; World Bank guidance links portability directly to lock-in risk and recommends assessing migration costs during procurement (World Bank, 2023, p. 24).

4. The Quantitative Model: Explanatory Variables, Interactions, and Hypotheses

Beyond compute and external exposure, the model includes AI governance capacity (*AIGC*), a sector-specific institutional capacity covering planning, procurement, contracting, coordination, risk management, and monitoring. Its narrower scope reduces overlap with general state capacity, although reverse causality remains possible.

AI governance capacity is denoted by:

$AIGC_i$

Data quality is denoted by:

DQ_i

DQ_i captures accuracy, completeness, timeliness, consistency, and interoperability.

Internal human and technical expertise (*HC*) remains part of *CSC* in the baseline model and is not entered again as a separate regressor. It may be analyzed independently only if *CSC* is reconstructed without *HC*, avoiding double counting.

The baseline relationship is:

$$SC_i = \beta_0 + \beta_1 AIC_i + \beta_2 AIGC_i + \beta_3 DQ_i + \beta_4 CSC_i + \beta_5 ED_i + \varepsilon_i$$

Here β_0 is the intercept and ε_i captures omitted influences on state capacity.

The baseline expects:

$$\beta_1 > 0, \beta_2 > 0, \beta_3 > 0, \beta_4 > 0$$

Thus *AIC*, *AIGC*, *DQ*, and *CSC* are expected, all else equal, to be positively associated with *SC*, while:

$$\beta_5 < 0$$

higher *ED* is expected to reduce *SC*. The article's central claim, however, concerns how *ED* changes the return to *AIC* rather than this direct effect alone.

Because the effect of dependence should vary with the level of AI use, the core specification adds an interaction term.

The interaction model is:

$$SC_i = \beta_0 + \beta_1 AIC_i + \beta_2 AIGC_i + \beta_3 DQ_i + \beta_4 CSC_i + \beta_5 ED_i + \beta_6 AIC_i ED_i + \varepsilon_i$$

The interaction term is:

$AIC_i ED_i$

It represents the expectation that additional usable computational capacity yields smaller gains when exposure is highly concentrated or critical.

The marginal effect of *AIC* is:

$$\partial SC_i / \partial AIC_i = \beta_1 + \beta_6 ED_i$$

The central hypothesis is $\beta_6 < 0$: higher *ED* reduces the positive marginal effect of *AIC*. Because *ED* measures concentrated or critical exposure rather than openness to external markets, empirical analysis should report marginal effects across *ED* levels rather than interpret the interaction coefficient in isolation.



For the linear model, the level of *ED* at which the marginal effect of *AIC* becomes *zero* is:

$$ED^* = -\beta_1 / \beta_6$$

If $ED < ED^*$, the marginal effect of *AIC* remains positive; above that threshold it may reach *zero* or become negative. ED^* is an algebraic result, not an empirical threshold, until β_1 and β_6 are estimated and the implied value falls within $[0,1]$.

The model deliberately distinguishes external reliance from technological dependence. Access to cloud services or external models can raise *AIC* when alternatives and governance safeguards are adequate; *ED* instead measures concentrated or critical exposure expected to constrain choice. *CSC* captures the state's capacity to manage that exposure. This separation avoids treating every external relationship as a sovereignty loss and leaves open the possibility of testing a separate, neutral measure of external reliance if future research examines nonlinear benefits and costs.

5. Empirical Testing Strategy, Validity, Causality, and Robustness

At this stage, the framework is a testable theoretical model, not an empirical result. The equations specify relationships to be confronted with data, while the composite indicators remain sensitive to indicator selection, normalization, weighting, and aggregation. Uncertainty and sensitivity analysis are therefore part of validation rather than optional additions (OECD & JRC, 2008, p. 118).

Empirical testing must also distinguish association from causation. More capable states may be better able to build sovereign computational capacity, so reverse causality is plausible. Panel data, temporal ordering, and country and year fixed effects can reduce some sources of confounding, but they do not by themselves identify causal effects. Causal claims require a design that credibly separates the effect of computational variables from pre-existing differences in state capacity.

Measurement must prevent structural overlap. Indicators used for *AIC* should not reappear in *CSC*, and *ED* should capture exposure and concentration rather than control dimensions already included in *CSC*. *CSG* cannot be entered in the same regression as *AIC* and *CSC* because $CSG = AIC - CSC$, which creates perfect linear dependence; it should be used descriptively or in an alternative specification. *CSC* may also mediate part of *ED*'s effect, while *AIGC* may itself reflect earlier state capacity, so total and direct effects must be distinguished explicitly.

The assumed causal graph should therefore be specified before estimation. Fixed effects can absorb stable unobserved heterogeneity but do not solve reverse causality or mediation; estimates should be described as causal only when supported by a credible identification strategy.

If SC is bounded in $[0,1]$, a linear model is useful as an interpretable baseline, but robustness should be assessed with a fractional-response model because *OLS* can predict outside the admissible range (Papke & Wooldridge, 1996, pp. 619–620). In nonlinear models, marginal effects must be derived from the link function, so $ED^* = -\beta_1/\beta_6$ applies only to the linear specification. Interaction tests should report expected marginal effects and confidence intervals across *ED* levels and use alternative operationalizations of external exposure to ensure results are not definition-specific.

Section Five: Examining the Model's Internal Consistency through Scenarios and Computational Simulation

This section uses scenarios and simulation only to test the model's internal consistency and sensitivity. It does not treat simulated outcomes as evidence about real states.

1. The Role of Simulation, Its Scope, and the Limits of Inference

Simulation is used here to trace the implications of the formal model under controlled assumptions, not to substitute for comparative empirical testing. Results are therefore conditional: they show what follows if specified values and coefficients hold. Because composite indicators depend on normalization, weighting, and aggregation choices, sensitivity analysis is necessary to determine



whether conclusions reflect the model's structure or arbitrary methodological decisions (OECD & JRC, 2008, p. 118).

All core variables are placed on a common $[0,1]$ scale for the exploratory simulations, where *zero* and one represent the defined lower and upper bounds. This makes unlike variables comparable within the model without implying that they share original units. Empirical applications must define and justify the actual bounds and normalization rules.

$$0 \leq AIC, CSC, AIGC, DQ, ED \leq 1$$

2. Constructing Scenarios and Hypothetical Values

For exploration, the theoretical model is instantiated with hypothetical coefficients whose signs match the proposed expectations. They are not empirical estimates and should not be interpreted as real-world effect sizes.

$$SC = 0.20 + 0.20 \cdot AIC + 0.20 \cdot AIGC + 0.15 \cdot DQ + 0.25 \cdot CSC - 0.10 \cdot ED - 0.20 \cdot AIC \cdot ED$$

The specification assigns positive effects to *AIC*, *AIGC*, *DQ*, and *CSC*, and negative effects to *ED* and the *AIC*×*ED* interaction.

The intercept of **0.20** keeps *SC* within $[0,1]$ for all admissible input values; under this hypothetical specification, *SC* ranges from **0.10** to **1.00**.

$$\partial SC / \partial AIC = 0.20 - 0.20 \cdot ED$$

Here $ED^* = 1$, so the marginal effect of *AIC* declines with dependence but reaches *zero* only at the upper bound; it does not become negative within $[0,1]$. Testing reversal would require a more negative β_6 or empirically estimated coefficients that place ED^* inside the interval.

3. Scenario Results and Comparison

Scenario 1 (high use, high control) sets *AIC* = **0.85**, *CSC* = **0.80**, *AIGC* = **0.80**, *DQ* = **0.75**, and *ED* = **0.20**, yielding *SC* $\approx$ **0.79** and *CSG* = **0.05**. It illustrates how high use can coexist with a small sovereignty gap when control is strong and exposure is limited; it is not a normative benchmark.

Scenario 2 (high use, low control) keeps *AIC* = **0.85**, *AIGC* = **0.80**, and *DQ* = **0.75** but sets *CSC* = **0.40** and *ED* = **0.75**, yielding *SC* $\approx$ **0.54** and *CSG* = **0.45**. The lower outcome is conditional on the hypothetical coefficients and illustrates the model's distinction between operational use and vulnerability in continuity and substitution.

Scenario 3 (moderate use, high control) sets *AIC* = **0.55**, *CSC* = **0.80**, *AIGC* = **0.80**, *DQ* = **0.75**, and *ED* = **0.20**, yielding *SC* $\approx$ **0.74** and *CSG* = **-0.25**. The negative gap denotes a relative control margin on the common scale, not a reserve of physical compute; the model still predicts higher *SC* if *AIC* rises while control and exposure remain favorable.

Scenario 4 (low use, low control) sets *AIC* = **0.35**, *CSC* = **0.30**, *AIGC* = **0.35**, *DQ* = **0.35**, and *ED* = **0.30**, yielding *SC* $\approx$ **0.42** and *CSG* = **0.05**. This shows why *CSG* cannot rank states by itself: a small gap may reflect either high use and control or low levels of both.

4. Sensitivity Analysis, Uncertainty, and Monte Carlo Simulation

Sensitivity analysis should vary weights, aggregation rules, coefficients, and variable values across plausible ranges. Monte Carlo simulation can extend this exercise by drawing from prespecified distributions, including values of β_6 near *zero* or with alternative signs. The purpose is to identify which conclusions are robust to assumptions, not to generate evidence for a hypothesis already encoded in the equation.

Because *AIC*, *CSC*, *AIGC*, *DQ*, and *ED* may be correlated across countries, independent draws can generate implausible combinations. Results should therefore be compared under independent sampling for global exploration and correlated sampling when empirical data permit estimation of the dependence structure.



Aggregation rules also require testing. A linear index assumes that strength on one dimension can compensate for weakness on another, an assumption that may be unsuitable when some sovereignty dimensions are critical. Linear aggregation should therefore be compared with geometric or threshold-based rules. As OECD/JRC guidance emphasizes, such sensitivity tests reveal how indicator selection, normalization, weighting, and aggregation shape results (OECD & JRC, 2008, p. 118). The goal is diagnostic robustness, not a rigid ranking of states.

5. Robustness, Limits of the Results, and the Transition to Comparative Testing

The simulations remain limited by hypothetical coefficients, possible bidirectional causality, and incomplete comparative data on auditability, contracts, and supplier exposure. Their value is therefore diagnostic: they make assumptions explicit, test internal consistency, and identify data needs. By separating concentrated external exposure from concrete capacities for audit, portability, substitution, continuity, and exit, the framework converts a broad claim about technological dependence into relationships that can be examined comparatively.

Section Six: Computational Sovereignty and State Capacity in the Global South

This section applies the framework to the Global South without treating the category as homogeneous or assuming that external reliance automatically produces dependence. It focuses on unequal access to compute, data, skills, and bargaining power, then considers how states can preserve autonomy of choice.

1. The Global South in the Global Distribution of Computational Capacity

The Global South is analytically important because countries occupy unequal positions in the digital value chain, not because external reliance is unique to them. As of June 2025, high-income countries accounted for **77%** of global colocation data-center capacity, compared with **18%** in upper-middle-income, **5%** in lower-middle-income, and less than **0.1%** in low-income countries; they also hosted **86%** of the top **500** high-performance computing systems and **97%** of their total capacity (World Bank, 2025, p. 39). These figures reveal a highly unequal material distribution of AI infrastructure.

Domestic infrastructure is not the only route to computational capacity because cloud services can substitute for local construction, though they create trade-offs involving latency, transfer costs, privacy, and data sovereignty (World Bank, 2025, p. 47). At the same time, commercial compute and cloud markets are geographically and commercially concentrated. This does not prove political dependence in every case, but it makes bargaining power, diversification, and switching capacity central to sovereignty analysis.

Countries within the same income group can occupy very different positions. The World Bank identifies middle-income regional computing hubs such as Brazil, India, Indonesia, Malaysia, South Africa, and Türkiye, supported by market size, connectivity, investment conditions, and domestic digital economies (World Bank, 2025, p. 46).

2. Inequalities in Data, Skills, and Sources of Dependence

Inequality also concerns data and contextual relevance. Global models may be poorly suited to public administration where local languages and materials are underrepresented; English remains overrepresented in web and open AI datasets (World Bank, 2025, p. 63). Yet this is not a fixed barrier: multilingual systems, open models, and local adaptation create room for states that cannot train frontier models. The relevant distinction is between capacity for frontier innovation and capacity for effective adaptation and application.

Skills form a third constraint. Infrastructure and data yield little administrative value without personnel able to use, evaluate, negotiate, and audit systems. Large cross-income gaps in basic and advanced digital skills therefore translate directly into differences in the ability to govern external technology (World Bank, 2025, p. 82).



3. External Reliance between Capacity Expansion and Deepening Vulnerability

These conditions shift the policy question from symbolic ownership to effective choice. A local data center dependent on one supplier for hardware, software, and maintenance may do little to reduce vulnerability, while diversified cloud contracts with clear portability rules and strong internal expertise may preserve more options. Domestic investment can still be strategically valuable when demand, power, connectivity, scale, and sustainability justify it, potentially lowering access costs, building skills, and strengthening bargaining power (World Bank, 2025, p. 39). The appropriate mix is therefore context-specific.

4. Strategies for Governing Dependence and Building Autonomy of Choice

Three complementary strategies follow: build selected domestic computing resources where scale, security, and continuity justify them; diversify external access through portability, open standards, and enforceable data and exit rights; and invest in public technical knowledge so institutions can assess vendors and switch systems when necessary. The objective is not complete computational independence, but strategic autonomy of choice—the ability to benefit from global technology markets without making public functions hostage to an irreplaceable vendor or technical layer.

5. The Computational Sovereignty Gap as a Framework for Comparison and Public Policy

The Computational Sovereignty Gap can help distinguish states with similar levels of AI use but different levels of institutional control. A state with audit capacity, portability rights, multiple vendors, and internal expertise differs materially from one with the same level of use but a single provider and weak technical knowledge. Cross-national comparison nevertheless faces data limitations, especially for contracts, exit capacity, and auditability, and may require combining international datasets with coded procurement documents and national strategies (World Bank, 2025, p. 47).

The framework is therefore better suited to comparative explanation than to normative ranking. Research designs should compare states with similar income or digitalization but different structures of computing access, procurement, and data governance, supplemented by contrasting cases that reveal mechanisms such as vendor negotiation and exit planning. The central question is not whether localization is maximized, but whether institutional room for maneuver expands or contracts.

Jansen et al. provide a useful precedent by measuring actual hosting locations and data-traffic routes rather than relying on sovereignty discourse alone (Jansen et al., 2023, p. 2). Computational sovereignty can be studied similarly through observed cloud providers, contractual terms, portability provisions, and deployed models. The unit of analysis may also vary: national indicators can be linked to ministry-, service-, or system-level data in multilevel designs because governance capacity may differ substantially within the same state.

This agenda extends computational political science by treating computational infrastructure itself as an object of political measurement and explanation. Compute is not only a methodological tool; it is a resource whose control is unevenly distributed across states, firms, and institutions.

Conclusion

The article has argued that technological adoption and state capacity should not be equated. Cloud services, external models, and specialized vendors can expand operational capacity while leaving the state with varying degrees of control over the conditions of use. Sovereign computational capacity captures this institutional control through auditability, portability, substitution, internal expertise, continuity, and exit, without equating sovereignty with technological self-sufficiency. The relevant criterion is the range of realistic options public institutions retain when a technological arrangement must be reviewed, modified, replaced, or terminated.

The Computational Sovereignty Gap formalizes the standardized difference between usable computational capacity and sovereign computational capacity, while the quantitative model separates the benefits of external access from the risks of concentrated or critical exposure. *AIC* captures usable



capacity, *ED* captures external exposure expected to reduce the marginal return to *AIC*, and *CSC* captures the state's ability to govern that exposure. The simulation scenarios are not empirical findings; they clarify the model's internal logic and the conditions that future comparative research should test.

The framework is especially relevant to the Global South, where access to compute, skills, and digital resources is highly unequal but does not mechanically determine dependence. States can trade for computing, adapt models, diversify suppliers, and build expertise. State capacity in the AI era therefore includes not only human, financial, and legal resources, but also the capacity to govern dependence on technical infrastructures located beyond the traditional boundaries of the public apparatus.

The central conclusion is that computational sovereignty is not the absence of dependence but the capacity to keep dependence governable, reviewable, and substitutable. A capable state is not necessarily the one with the largest domestic compute base or the most AI applications, but one that can use such resources while preserving meaningful institutional choice. Measuring the gap between use and control provides a way to analyze how the conditions of public capacity are being redistributed across digital infrastructures, firms, and jurisdictions.



References

1. Busuioc, Madalina. "Accountable Artificial Intelligence: Holding Algorithms to Account." *Public Administration Review*, vol. 81, no. 5, 2021. <https://doi.org/10.1111/puar.13293>
2. Couture, Stéphane & Sophie Toupin. "What Does the Notion of 'Sovereignty' Mean When Referring to the Digital?" *New Media & Society*, vol. 21, no. 10, 2019, pp. 2305–2322. <https://doi.org/10.1177/1461444819865984>
3. Estima, Walmir. "Popular Digital Sovereignty and Infrastructural Dependencies: The Contrate Quem Luta Project." *Digital Society*, vol. 5, Article 27, 2026. <https://doi.org/10.1007/s44206-026-00266-2>
4. Farrell, Henry & Abraham L. Newman. "Weaponized Interdependence: How Global Economic Networks Shape State Coercion." *International Security*, vol. 44, no. 1, 2019.
5. Fukuyama, Francis. "What Is Governance?" *Governance*, vol. 26, no. 3, 2013.
6. Hickok, Merve. "Public Procurement of Artificial Intelligence Systems: New Risks and Future Proofing." *AI & Society*, vol. 39, no. 3, 2024. <https://doi.org/10.1007/s00146-022-01572-2>
7. Jansen, Bernardus, Natalia Kadenko, Dennis Broeders, Michel van Eeten, Kevin Borgolte & Tobias Fiebig. "Pushing Boundaries: An Empirical View on the Digital Sovereignty of Six Governments in the Midst of Geopolitical Tensions." *Government Information Quarterly*, vol. 40, no. 4, Article 101862, 2023. <https://doi.org/10.1016/j.giq.2023.101862>
8. OECD & Joint Research Centre (JRC). *Handbook on Constructing Composite Indicators: Methodology and User Guide*. Paris: OECD Publishing, 2008.
9. OECD. *Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions*. Paris: OECD Publishing, 2025. <https://doi.org/10.1787/795de142-en>
10. Papke, Leslie E. & Jeffrey M. Wooldridge. "Econometric Methods for Fractional Response Variables with an Application to 401(k) Plan Participation Rates." *Journal of Applied Econometrics*, vol. 11, no. 6, 1996, pp. 619–632.
11. Roberts, Huw. "Digital Sovereignty and Artificial Intelligence: A Normative Approach." *Ethics and Information Technology*, vol. 26, Article 70, 2024. <https://doi.org/10.1007/s10676-024-09810-5>
12. World Bank. *Digital Progress and Trends Report 2025: Strengthening AI Foundations*. Washington, DC: World Bank, 2025.
13. World Bank. *Institutional and Procurement Practice Note on Cloud Computing: Cloud Assessment Framework and Evaluation Methodology*. Washington, DC: World Bank, 2023.