Digital sovereignty and AI governance: from claim to instrument in a comparison of the Brazilian AI Plan and America's AI Action Plan

Darci de Borba Santos Júnior

https://doi.org/10.1590/SciELOPreprints.17393

Submitted on: 2026-08-11
Posted on: 2026-08-17 (version 1)
(YYYY-MM-DD)

SOBERANIA DIGITAL E GOVERNANÇA DA INTELIGÊNCIA ARTIFICIAL: DA REIVINDICAÇÃO AO INSTRUMENTO NA COMPARAÇÃO ENTRE O PLANO BRASILEIRO DE IA E O AMERICA'S AI ACTION PLAN

DARCI DE BORBA SANTOS JÚNIOR
ORCID: https://orcid.org/0000-0001-9851-5034
<darci17@gmail.com>
Instituto de Pesquisa Econômica Aplicada (Ipea). Brasília, Distrito Federal (DF), Brasil

RESUMO: A soberania digital tornou-se vocabulário comum das estratégias nacionais de inteligência artificial (IA), mas a literatura descreve reivindicações de soberania sem explicar por que um mesmo vocabulário produz configurações de política opostas. Persiste, assim, uma lacuna teórica: falta um quadro que conecte a camada da reivindicação à camada do instrumento e que acomode a soberania exercida sobre o acesso de terceiros. O artigo compara o Plano Brasileiro de Inteligência Artificial (PBIA 2024-2028) e o America's AI Action Plan (AAP, 2025) para responder como cada plano traduz reivindicações de soberania em instrumentos e o que essa divergência revela sobre a relação entre posição estrutural na produção de IA e estratégia de soberania disponível. Adota-se realismo crítico e lógica retrodutiva, com análise documental comparativa de dois casos de máxima diferença, codificados em cinco dimensões e em três camadas: soberania declarada, instrumentada e exercida. Os achados mostram que o PBIA declara soberania de modo transversal, mas a instrumenta por crédito subsidiado, destinando 0,45% do envelope ao eixo de governança e cerca de 2% do investimento estrutural ao habilitador de energia, ao passo que o AAP instrumenta a soberania por desregulamentação, aceleração de licenciamento, controle de exportações e exportação da pilha tecnológica integrada. A contribuição teórica está na distinção entre as três camadas, na extensão da tipologia de papéis do Estado para incorporar o controle sobre o acesso alheio e em três

proposições com condições de contorno explicitadas. A análise documental capta desenho institucional, não dinâmicas de implementação.

Palavras-chave: soberania digital, governança da inteligência artificial, instrumentos de política pública, capacidade estatal, análise comparativa

DIGITAL SOVEREIGNTY AND AI GOVERNANCE: FROM CLAIM TO INSTRUMENT IN A COMPARISON OF THE BRAZILIAN AI PLAN AND AMERICA'S AI ACTION PLAN

ABSTRACT: Digital sovereignty has become a common vocabulary of national artificial intelligence (AI) strategies, yet the literature describes sovereignty claims without explaining why the same vocabulary yields opposite policy configurations. A theoretical gap remains: no framework connects the layer of the claim to the layer of the instrument, and none accommodates sovereignty exercised over the access of third parties. This article compares the Brazilian Artificial Intelligence Plan (PBIA 2024-2028) and America's AI Action Plan (AAP, 2025) in order to answer how each plan translates sovereignty claims into policy instruments, and what that divergence reveals about the relation between structural position in AI production and the sovereignty strategy available to a state. The study adopts critical realism and retroductive inference, with comparative documentary analysis of two maximally different cases coded along five dimensions and three layers: declared, instrumented and exercised sovereignty. The findings show that the PBIA declares sovereignty across all of its axes yet instruments it through subsidised credit, allocating 0.45 per cent of the envelope to the governance axis and around 2 per cent of structural investment to the electricity enabler, whereas the AAP instruments sovereignty through deregulation, permitting acceleration, export controls and the export of an integrated technology stack. The theoretical contribution lies in distinguishing the three layers, in extending

the typology of state roles to incorporate control over the access of others, and in three propositions with stated boundary conditions. Documentary analysis captures institutional design rather than implementation dynamics.

Keywords: digital sovereignty, artificial intelligence governance, policy instruments, state capacity, comparative analysis

INTRODUCTION

National artificial intelligence strategies have become the standard instrument through which governments state their position on a technology whose economic and security implications remain unsettled. More than sixty countries have published such documents, and the pace accelerated after generative models entered general use (Papyshev & Yarime, 2023; Radu, 2021). These texts allocate resources, create oversight bodies and define what counts as an acceptable risk. Bareis and Katzenbach (2022) show that they also perform a narrative function, since the description of an inevitable technological future helps produce the conditions of its realisation.

One term recurs across almost all of them. Digital sovereignty migrated from a specialised concern of European industrial policy to a general frame through which states justify intervention in the technology sector (Pohle & Thiel, 2020). Its usefulness lies in its breadth. The same expression can designate the protection of citizens' data from foreign surveillance, the construction of domestic computing capacity, the regulation of transnational platforms, the control of semiconductor supply chains or the assertion of jurisdiction over algorithmic decisions. That breadth has made the term politically indispensable and analytically unstable (Couture & Toupin, 2019).

Research has responded to this instability along four separate lines. A first line has mapped the meanings of sovereignty claims and the actors that voice them (Couture & Toupin, 2019; Pohle & Thiel, 2020), and has proposed normative criteria for distinguishing control from legitimate

authority (Floridi, 2020; Roberts, 2024). A second line has inventoried the material requirements of autonomy, most explicitly in the enabler framework of Belli (2023), which decomposes AI sovereignty into data, algorithms, computing capacity, connectivity, electricity, education, cybersecurity and regulation. A third line has produced typologies of the roles states assume in AI governance, of which the development, control and promotion triad of Papyshev and Yarime (2023) is the most widely used, alongside Radu's (2021) account of hybrid governance and functional indetermination and the comparative analysis of state responsibility by Djeffal et al. (2022). A fourth line, more recent, has read state responses to AI concentration as a process of international reordering, in which export controls and infrastructure partnerships produce techno-blocs (Weymouth, 2025).

Each line explains something the others do not, and their juxtaposition exposes what none of them addresses. The first line describes claims but does not connect them to the instruments through which claims are pursued, so that a plan declaring sovereignty in every section and a plan declaring it in one are treated as comparable expressions of the same phenomenon. The second line inventories requirements but does not explain why a given state addresses some enablers and not others. The third line was constructed from corpora of strategies published between 2017 and 2021, before export restriction became a principal instrument of AI policy, and its categories therefore describe functions the state performs at home; sovereignty exercised by conditioning what other states may acquire has no place in the typology. The fourth line captures that outward movement but operates at the level of the international system, without returning to the domestic policy configurations that make it possible.

The gap is theoretical and conceptual rather than empirical. What is missing is an account that links the layer at which sovereignty is claimed to the layer at which it is instrumented, and that specifies why the relation

between the two varies with a state's position in the international division of AI production. The absence has consequences for how comparative work reads its material. Declared ambition is read as strategy, so that a plan announcing an ample envelope appears more sovereign than one announcing none. The absence of a declared accountability agenda is read as a governance vacuum, when it may be the trace of an active redirection of governance capacity. Assessments of national plans consequently oscillate between crediting intentions and dismissing rhetoric, without an intermediate category in which the alignment between claim and instrument can be examined.

The problem that follows is scientific and investigable. If sovereignty claims are ubiquitous and instruments are heterogeneous, comparison requires separating the two and asking what accounts for their alignment or misalignment. Uncertainty about this relation matters for three reasons. It affects how the effectiveness of national AI policy is assessed, since a plan may be coherent in its claims and incoherent in its instruments. It affects the transferability of policy prescriptions, since instruments that operate in one structural position may be unavailable or counterproductive in another. And it affects the interpretation of governance asymmetries between countries that produce and countries that adopt AI technologies, an asymmetry that has become more pronounced as access to computing capacity has been made conditional.

The article therefore asks: how do the Brazilian AI Plan and America's AI Action Plan differ in the way they translate digital sovereignty claims into policy instruments, and what does that divergence reveal about the relation between a state's position in the international division of AI production and the sovereignty strategy available to it?

The general objective is to explain the divergence between the two plans by reconstructing, for each case, the translation from sovereignty claim to policy instrument. Three specific objectives organise the analysis.

The first is to characterise the sovereignty claims declared in each plan across five analytical dimensions. The second is to identify the instruments, allocations and institutions through which each plan operationalises those claims, and to establish where declaration and instrumentation diverge. The third is to derive, from the observed divergence, propositions about the relation between structural position and sovereignty strategy, together with the conditions under which those propositions are expected to hold.

The study adopts a critical realist position and a retroductive logic of inquiry, applied through comparative documentary analysis of two maximally different cases. Brazil and the United States occupy opposite positions in the AI value chain while issuing wide-ranging national plans within a twelve-month interval, which holds the technological moment and the genre of the document approximately constant. Coding proceeds along five dimensions, namely technological autonomy, governance structure, ethical accountability, funding and infrastructure, and international positioning, and along three layers, namely declared, instrumented and exercised sovereignty.

Three contributions are anticipated. The first is conceptual: the separation of declared, instrumented and exercised sovereignty gives the literature an intermediate category between rhetoric and capability, and makes the translation gap itself an object of analysis rather than a residual. The second is theoretical: the comparison supports an extension of the typology of state roles in AI governance, adding an outward axis in which leverage derives from conditioning the access of other states, and it yields three propositions with explicit boundary conditions. The third is methodological: the coding protocol combines textual indicators of claim with allocative indicators of instrument, which allows the divergence between the two to be observed rather than inferred.

The article proceeds as follows. Section 1 reconstructs the theoretical debate and presents the analytical framework. Section 2 sets out the

philosophical positioning, the design and the coding procedure. Section 3 reports the results for each case and the cross-case comparison. Section 4 discusses the findings, formulates the propositions and examines alternative explanations. Section 5 states the contributions. The final section answers the research question, delimits the scope of the inferences and proposes a research agenda.

1. THEORETICAL BACKGROUND

1.1 Digital sovereignty: from claim to capability

Classical sovereignty designates supreme authority over a delimited territory, combining internal supremacy with external independence. Digital environments unsettle both attributes, since data flows ignore borders, the infrastructure carrying them is largely privately held, and the standards governing interoperability are set in technical fora rather than legislatures. The expectation that these features would render territorial authority obsolete did not survive the following decade.

Pohle and Thiel (2020) attribute the return of sovereignty discourse to the weakening of two earlier paradigms, cyber-libertarianism and multi-stakeholderism, and characterise what replaced them less as a legal category than as a discursive practice through which heterogeneous actors articulate claims to control. They distinguish three recurrent strands: state autonomy over infrastructure and data flows, economic autonomy in relation to foreign providers, and individual self-determination. Couture and Toupin (2019) arrive at a convergent conclusion by a different route, showing that states, firms, social movements and indigenous peoples invoke sovereignty with meanings that are frequently antagonistic. Their analysis warns against treating the term as self-evident, a warning that applies with particular force to comparative work.

Two responses to that instability have been proposed. Floridi (2020) reframes the question as one of the relation between control and legitimacy, reading the European project as an attempt to reassert public authority over infrastructures that have acquired quasi-governmental effects. Roberts (2024) develops the point systematically, arguing that the prevailing descriptive conception, which asks who exercises control, cannot distinguish legitimate from illegitimate control, and proposing a normative reconstruction centred on authority. Applying it to large technology firms, he concludes that their high degree of control coexists with weak input legitimacy and questionable output legitimacy, which justifies treating them as quasi-sovereigns rather than sovereigns. The distinction bears directly on national AI plans, which are among other things attempts to reconvert corporate control into public authority.

A separate strand has asked whether sovereignty claims are matched by capability. Calderaro and Blumfelde (2022) examine the European security agenda and conclude that, without a leading domestic industry and a coherent defence strategy, the Union possesses few instruments beyond regulation, which they describe as a false promise of sovereignty. Mügge (2024) asks for whom and to whose benefit European AI sovereignty is pursued and finds the answer to vary across member states and industrial constituencies. Papyshev and Chan (2026) formalise the underlying dilemma, showing that stringent regulation operates as a sovereignty strategy only when it constrains foreign firms disproportionately and when the foreign technological advantage is material without being overwhelming. Ashraf and Veneziano (2026) review the emerging literature on AI sovereignty and identify infrastructural dependence as the condition that most often separates strategic agency from its rhetorical assertion.

The material substrate of these debates has been documented with increasing precision. Van der Vlist et al. (2024) trace the convergence of AI and a small number of cloud providers, showing that the industrialisation of

AI has produced dependencies spanning infrastructure, models and applications, and that the same firms have integrated all three layers into their commercial offerings. Weymouth (2025) reads state responses as digital disintegration, in which export controls, infrastructure partnerships and platform restrictions produce coalitions that coordinate control over critical inputs while excluding rivals. Sovereignty in this account is not a residue of the territorial state but an emergent form of international ordering.

For countries outside the technological core the configuration differs. Belli (2023) proposes the Key AI Sovereignty Enablers framework, identifying eight constitutive components of a sovereignty stack and providing an operational vocabulary for assessing whether a strategy addresses the full set of dependencies. Schiavi et al. (2025) analyse four Brazilian policy documents on digitisation and find that their content does not sustain a development agenda capable of fostering sovereignty, warning that the gap is likely to deepen dependence. Sengupta et al. (2025) offer a qualification by examining the interaction between AI and digital public infrastructure in India and Brazil, suggesting that publicly controlled infrastructural layers may constitute a distinct pathway to autonomy for middle-income countries.

1.2 State roles in AI governance and the limits of existing typologies

The literature on what states do with AI policy has produced categories that this article both uses and revises. Papyshev and Yarime (2023) analyse thirty-one national strategies through content analysis and topic modelling, identify thirteen state functions and aggregate them into three roles. Development describes states that build capability the market has not produced; control describes states that regulate; promotion describes states that remove obstacles to private capability formation. Their finding that post-Soviet and East Asian countries prioritise development, European countries control, and the United States, the United Kingdom and

Ireland promotion, has become a reference point for comparative work. Djeffal et al. (2022) reach compatible conclusions from an examination of how conceptions of state responsibility structure instrument selection.

Radu (2021) shifts attention from roles to modes, showing that strategies released between 2017 and 2019 exhibit a predominance of ethics-oriented over rule-based arrangements and a preference for functional indetermination, understood as the deliberate maintenance of ambiguity concerning obligations and enforcement. Ulnicane et al. (2021) document the oscillation of policy documents between promissory and precautionary registers. Corrêa et al. (2023), reviewing two hundred governance instruments, identify seventeen recurrent principles and a persistent distance between declaration and implementation, a distance that Birkstedt et al. (2023) confirm in their systematic review, where limited understanding of implementation appears as the central knowledge gap.

The corpora on which these typologies rest were assembled before 2022. Export restriction, compute conditionality and the packaging of national technology stacks for export were not yet the principal instruments of AI policy, and the categories accordingly describe functions performed within the state's own jurisdiction. When a state makes the acquisition of computing capacity by third countries conditional on alignment, it exercises control, but not over its own market; it promotes its industry, but by restricting the options of others. The triad registers this as promotion, which loses the distinctive feature of the instrument. The limitation is one of period rather than of construction, and it is the point at which the present study intervenes.

Evidence from public administration research supplies the second element of the framework. Busuioc (2021) argues that algorithmic systems generate accountability deficits that existing administrative controls address poorly, and specifies explainability, contestability and the availability of remedies as the functions that must be secured. Dunleavy and Margetts

(2025) situate these developments within a third wave of digital era governance in which data-intensive technologies reshape information regimes and the allocation of state functions. Chandra and Feng (2026) map a decade of scholarship and find that algorithmic decision-making dominates the agenda while implementation and contextual conditions remain underexamined. For Brazil, Ribeiro and Segatto (2025) show through survey microdata that AI adoption by public agencies remains restricted to a small number of organisations and is strongly associated with pre-existing information technology capacity, with pronounced heterogeneity across federal and subnational entities. Filgueiras (2022) argues that national AI strategies constitute distinct policy regimes whose politics cannot be reduced to technical content.

Comparative work on the great powers completes the background. Hine and Floridi (2024) trace the divergence between American and Chinese approaches to differences in philosophy of technology rather than to momentary geopolitical alignment. Roberts et al. (2021) document the Chinese configuration of policy, ethics and regulation, and Hung (2025) shows that cyber sovereignty in Chinese official discourse extends a state-centric conception of international order into the digital domain. Ishkhanyan (2025) formulates the resulting tension as a sovereignty-internationalism paradox. Santos Júnior and Brinkhues (2026) assess America's AI Action Plan against the regulatory practices identified in the peer-reviewed literature and find selective convergence in cybersecurity, technological sovereignty and standardisation alongside divergence in rights protection and environmental responsibility, proposing a reading of AI governance as a multidimensional and adaptive regime. Their assessment is normative in orientation, since the plan is measured against a standard external to it. The present article poses a different question, structural rather than normative, by asking what the configuration of instruments reveals about the position from which each state acts.

1.3 Analytical framework

Three propositions extracted from the literature organise the analysis. Sovereignty is composite rather than unitary, being distributed across policy domains with different capital intensities and different degrees of international contestation, which implies that comparison must proceed dimension by dimension. Claims and capabilities are analytically separable, so that the presence of sovereignty vocabulary in a document establishes nothing about the state's capacity to exercise control. And the meaning of sovereignty is conditioned by position in the international division of technological labour, so that identical vocabulary may encode opposite strategic problems.

The framework operationalises these propositions through three layers. Declared sovereignty comprises the claims inscribed in the policy text, including stated premises, objectives and normative commitments. Instrumented sovereignty comprises the instruments actually selected and funded, including budgetary allocations, legal acts, institutions created and conditionalities imposed. Exercised sovereignty comprises the capability effectively deployed, observable in execution data, in institutions that have begun to operate and in control over specific layers of the stack. The distance between the first and the second layer is designated the translation gap; the distance between the second and the third, the execution gap. Figure 1 represents the framework.

Figure 1
Three-Layer Analytical Framework for the Comparative Analysis of Digital Sovereignty in National AI Plans

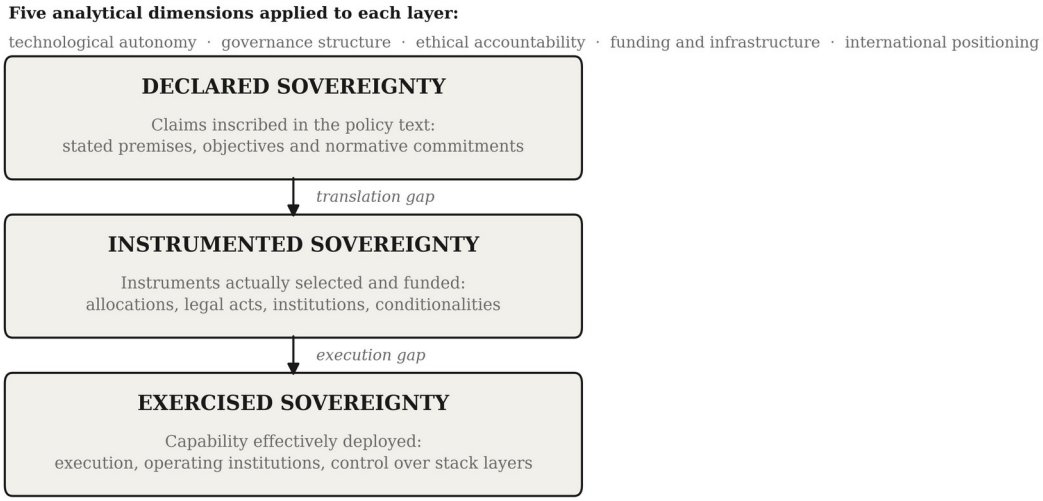

Note. Source: prepared by the author.

Five dimensions are applied to each layer. Technological autonomy concerns control over the constitutive layers of the AI stack, following the enabler logic of Belli (2023). Governance structure concerns the institutional architecture for coordination, monitoring and decision, drawing on Papyshev and Yarime (2023) and Radu (2021). Ethical accountability concerns the mechanisms through which AI systems are rendered answerable, following Busuioc (2021) and Corrêa et al. (2023). Funding and infrastructure concerns the volume, composition and instruments of investment and the material commitments in compute and energy, following van der Vlist et al. (2024). International positioning concerns how the plan situates the country in the global technological order, following Weymouth (2025) and Ishkhanyan (2025). Table 1 presents the operationalisation.

Table 1

Operationalisation of the Five Analytical Dimensions Across the Declaration and Instrumentation Layers

Dimension	Definition	Indicators of declaration	Indicators of instrumentation	Reference
Technological autonomy	Degree of domestic control sought over the layers of the AI stack	Statements on sovereignty, autonomy, national capability	Actions and allocations by enabler; procurement of compute; model and semiconductor programmes	Belli (2023)
Governance structure	Institutional architecture for coordination, monitoring and decision	Bodies foreseen; stated coordination mandate	Bodies effectively created; legal instruments; reporting obligations; enforcement mechanisms	Papyshev and Yarime (2023); Radu (2021)
Ethical accountability	Mechanisms rendering AI systems answerable	Principles; foreseen oversight bodies; commitments to transparency	Allocation to the governance function; binding rules; status of the regulatory framework	Busuioc (2021); Corrêa et al. (2023)
Funding and infrastructure	Volume, composition and type of investment instrument	Announced envelope; stated infrastructure objectives	Executed or secured resources; instrument type (credit, grant, regulatory, fiscal); energy and permitting provisions	Van der Vlist et al. (2024)
International positioning	Placement of the country in the global technological order	Statements on cooperation, leadership, multilateralism	Export controls; export programmes; alliance instruments; participation in governance fora	Weymouth (2025); Ishkhanyan (2025)

Note. Source: prepared by the author.

2. METHOD

The study is conducted from a critical realist position (Bhaskar, 2008; Sayer, 2000). Ontologically, it treats the international division of AI production as a structure that exists independently of the policy documents analysed and that conditions, without determining, the instruments available to each state. Policy texts are read as products of that structure and as interventions upon it. Epistemologically, the position implies that observable policy configurations are the empirical trace of mechanisms that are not directly observable, so that the analytical task is to move from the configuration to the mechanism that would render it intelligible. Axiologically, the study recognises that sovereignty is a contested political category and that its invocation serves interests, which requires that the analysis avoid adopting the normative standpoint of either document.

The logic of inquiry is retroductive (Danermark et al., 2019; Fletcher, 2017). Neither deduction nor induction fits the task, since the study does not test hypotheses derived from an established theory and does not generalise from a sample. It asks what must be the case for the observed divergence between two policy configurations to be possible, and formulates the answer as propositions available for subsequent examination. Critical realist case research of this kind supports explanatory rather than statistical inference, and its claims are conditional on the structures posited (Wynn & Williams, 2012). The propositions advanced in Section 4 are stated with the boundary conditions under which they are expected to hold, in accordance with the requirement that a theoretical contribution specify its own limits (Whetten, 1989).

The design is a comparative analysis of two cases based on documentary evidence (Bowen, 2009). Documentary analysis is appropriate because the object of inquiry is the conception of sovereignty inscribed in official policy, which is accessible through the texts, the institutional architecture they establish and the allocative decisions they record.

Case selection follows a most-different-systems logic. Both documents are comprehensive national AI plans issued by the executive branch within a twelve-month interval, which holds constant the technological moment, the genre conventions of the strategy document and the general policy problem addressed. They differ systematically in the variables of theoretical interest: position in the AI value chain, fiscal capacity, regulatory tradition and role in the international system. Where a common outcome, in this case the invocation of sovereignty, appears under divergent structural conditions, the comparison isolates the conditions under which the outcome takes different forms (Eisenhardt, 1989). Two cases do not support statistical generalisation and are not intended to. They support analytical generalisation, meaning the transfer of an explanation to cases sharing the relevant structural features.

The unit of analysis is the national AI plan understood as a configuration of sovereignty claims and policy instruments. The unit of observation is the thematic segment of the document, defined as the passage corresponding to an axis, a pillar, an action or an instrument. The temporal horizon is the period between the publication of the Brazilian plan in July 2024 and August 2026, which covers the first full year of execution of both documents.

The primary corpus comprises two documents. The first is the Brazilian Artificial Intelligence Plan 2024-2028 (Plano Brasileiro de Inteligência Artificial, PBIA), launched on 29 July 2024 at the plenary session of the National Council for Science and Technology and revised in a final version released in June 2025. The second is America's AI Action Plan, published on 23 July 2025 by the Office of Science and Technology Policy in compliance with Executive Order 14179.

A secondary corpus was assembled to establish the instrumentation and execution layers, which the primary documents do not by themselves record. For Brazil it comprises the resolution creating the plan's working

group, published on 12 May 2025, the legislative record of Bill 2338/2023, the monitoring materials of the Brazilian Artificial Intelligence Observatory coordinated by NIC.br, and official implementation reports issued during 2025. For the United States it comprises Executive Orders 14318, 14319 and 14320, signed concurrently with the plan, and the executive order of December 2025 concerning state AI legislation.

Inclusion in the secondary corpus required that a document either create an obligation, allocate a resource or report execution. Commentary, press coverage and advocacy material were consulted to locate primary sources and to identify contested points, but were not treated as evidence of instrumentation.

Analysis proceeded in four stages. Each document was first read integrally and segmented according to its own structure of axes or pillars. Segments were then coded against the five dimensions and assigned to the declaration or instrumentation layer, with the rule that a segment was assigned to declaration when it stated an objective, a premise or a principle, and to instrumentation when it allocated a resource, created a body, imposed an obligation or established a conditionality. A third stage constructed dimension-level descriptions for each case and compared them through a synthesis matrix registering the dominant orientation, the principal instruments and the observed limitations. A fourth stage moved retroductively from the compared configurations to the mechanisms capable of accounting for them, and formulated the propositions reported in Section 4.

Where a document declares an objective without an accompanying instrument, or funds an instrument without a corresponding declaration, the divergence was recorded as such. This procedure makes the translation gap an observable feature of the material rather than an interpretive judgement applied to it.

Three procedures were adopted. Financial and institutional data were triangulated across the policy text, official implementation reports and independent monitoring sources, and divergences between sources are reported rather than resolved by selection. Two points remain genuinely contested in the available material and are treated accordingly: the disaggregation of the fifty-four actions of the Brazilian plan, which secondary sources describe inconsistently, and the count of federal policy actions in the American plan, stated as more than ninety in the document and as one hundred and three in an independent enumeration. Claims about content were kept separate from claims about execution throughout.

The design carries limits that bound the conclusions. Documentary analysis captures declared intentions and institutional design, not implementation dynamics, bureaucratic resistance or the informal arrangements through which policy is enacted. The absence of interviews precludes access to the deliberative processes that produced the documents. Both plans remain in execution, so assessments of the exercised layer are provisional. The inferences drawn are explanatory and conditional, and no causal claim is made about the effect of either plan on technological outcomes.

3. RESULTS

3.1 The Brazilian Artificial Intelligence Plan

Declared sovereignty in the Brazilian plan is transversal. The document states technological and data sovereignty among its premises, commits to the domestic development of AI processors, justifies a sovereign cloud infrastructure by the need to secure national technological autonomy, and links the reduction of algorithmic bias to data sovereignty. The subtitle, AI for the Good of All, situates autonomy as instrumental to social inclusion rather than to competitive advantage. Sovereignty appears in the plan not

as one objective among others but as the frame within which the remaining objectives are stated.

Instrumented sovereignty presents a different distribution. The announced envelope of R$ 23.03 billion for 2024-2028 is allocated across five axes. AI infrastructure and development receives R$ 5.79 billion across thirteen actions; diffusion, training and skills in AI receives R$ 1.15 billion across eight; AI for the improvement of public services receives R$ 1.76 billion across nineteen; AI for business innovation receives R$ 13.79 billion across nine; and support to the regulatory and governance process receives R$ 103.25 million across five. A set of immediate-impact actions with a twelve-month horizon receives R$ 435.04 million. Close to sixty per cent of the announced total is allocated to business innovation, and approximately 0.45 per cent to the axis dedicated to regulation and governance.

The distribution by sovereignty enabler is more revealing than the distribution by axis. Johansson Neto et al. (2024) map the plan's structural allocations onto the framework of Belli (2023) and find that six of the eight enablers are addressed, with close to half of structural investment directed to algorithms, approximately twenty per cent each to data and to computing capacity, eleven per cent to education, two per cent to electricity and one per cent to cybersecurity. Connectivity and regulation receive no dedicated structural allocation. The plan therefore concentrates on the layers where domestic scientific capacity already exists and is comparatively silent on the layers where dependence is most acute and where capital requirements are highest.

Material commitments are concrete without being large in international comparison. The plan foresees the acquisition of a supercomputer among the world's five most powerful, and the upgrade of the Santos Dumont installation at the National Laboratory for Scientific Computing has been reported as completed. A sovereign cloud initiative of the order of R$ 1 billion, channelled predominantly through Serpro,

addresses public-sector data residency. The development of advanced language models in Portuguese trained on national data constitutes the model-layer commitment. Human capital objectives include the training of twenty thousand professionals annually until 2028, fifty chip designers and one hundred infrastructure technicians. Seven National Institutes of Science and Technology dedicated to AI were selected during the first year.

The governance architecture is interministerial. The plan text foresees a superior council, an executive committee and five thematic chambers. The instance that became operational is the working group instituted on 12 May 2025 for a four-year term, composed of fifteen bodies and entities coordinated by the Ministry of Science, Technology and Innovation, including the Office of the Chief of Staff of the Presidency, the Ministry of Management and Innovation in Public Services, the Ministry of Foreign Affairs and the financing agencies Finep, BNDES, Capes, CNPq and Embrapii. The group monitors execution, proposes adjustments to the Interministerial Committee for Digital Transformation and submits an annual report. Monitoring is externalised to the Brazilian Artificial Intelligence Observatory (OBIA), coordinated by NIC.br.

The regulatory function sits outside the plan. Bill 2338/2023 was approved by the Senate on 10 December 2024 after the consideration of more than two hundred amendments and transmitted to the Chamber of Deputies in March 2025, where it remained, in mid-2026, awaiting the opinion of the rapporteur of the special committee constituted to examine it, with successive postponements of the anticipated vote. Civil society organisations have contested the text under discussion on grounds including the breadth of exceptions for facial recognition by security forces, the removal of labour safeguards and the treatment of copyright in training data. The National Centre for Algorithmic Transparency and Trustworthy AI foreseen under the fifth axis, together with a multisectoral ethics committee

and a regulatory framework projected for 2027, had not been established within the period analysed.

Exercised sovereignty can be assessed only partially. By December 2025, official reporting indicated that twenty-five of the fifty-four actions had produced deliverables and fifteen had been initiated, with R$ 6.47 billion secured of the announced R$ 23.03 billion, approximately R$ 4.5 billion through Finep and R$ 1.2 billion through BNDES. The predominance of these agencies confirms that the financial architecture rests on credit and subsidised financing rather than on direct budgetary appropriation, a feature that analysts at the Institute of Economics of the University of Campinas identify as poorly suited to frontier innovation and compound with the absence of public procurement for innovation. The same assessment notes an internal tension in the sovereign cloud proposal, which contemplates contracting foreign cloud service providers, and observes that intellectual property and open-source model strategies are not addressed.

3.2 America's AI Action Plan

Declared sovereignty in the American plan is expansive rather than defensive. The document articulates the objective of securing unquestioned and unchallenged global technological dominance and organises its measures around winning a race whose prize is the capacity to set global standards. The threat it identifies is not dependence on foreign technology but the possibility that another power might determine the terms on which AI is developed and diffused. Autonomy is presupposed rather than pursued.

The plan is organised in three pillars and states that it contains more than ninety federal policy actions, a figure that an independent enumeration by Stanford's Institute for Human-Centered Artificial Intelligence raises to one hundred and three. Its doctrinal premise was established before its publication. Executive Order 14110 of October 2023, which had governed the safe and trustworthy development of AI, was revoked on the first day of

the administration and replaced by Executive Order 14179 of 23 January 2025, which commissioned the plan itself and framed the preceding arrangement as an onerous regulatory regime.

The first pillar addresses domestic conditions for innovation. Agencies are directed to review and rescind regulations identified as impeding AI development; the Office of Management and Budget is instructed to consider states' regulatory climates in the allocation of discretionary funds; the NIST AI Risk Management Framework is to be revised to remove references to misinformation, diversity, equity and inclusion, and climate change. The pillar supports open-source and open-weight models, expands compute access through the National Artificial Intelligence Research Resource, promotes regulatory sandboxes and articulates a workforce agenda encompassing a research hub, retraining and AI literacy. Federally funded scientific data is treated as a strategic asset to be made available for model development.

The second pillar addresses material conditions. It provides categorical exclusions under the National Environmental Policy Act for data centre construction, expands the FAST-41 permitting regime, contemplates the use of federal lands and directs attention to grid modernisation with emphasis on dispatchable sources including advanced geothermal and nuclear generation. It seeks the restoration of domestic semiconductor manufacturing, including simplification of requirements administered by the CHIPS Program Office, and provides for high-security data centres serving defence and intelligence functions.

The third pillar addresses external projection. Its central instrument is the export of the full American technology stack, comprising hardware, models, software, applications and standards, to allied countries. The pillar seeks to counter competing influence in international governance fora, tightens semiconductor export controls through location verification, diversion monitoring and the closure of gaps in manufacturing equipment

through subsystem regulation, and establishes plurilateral technology alliances. Biosecurity measures concerning nucleic acid synthesis screening are included.

Instrumentation followed rapidly. Three executive orders were signed with the plan. Executive Order 14318 accelerates federal permitting of data centre infrastructure and offers federal financial support to projects meeting thresholds of investment of at least five hundred million dollars or electrical load above one hundred megawatts. Executive Order 14319 establishes principles of ideological neutrality and truth-seeking as conditions of federal procurement, with implementing guidance from the Office of Management and Budget and mandatory amendment of existing contracts. Executive Order 14320 creates the American AI Exports Program, operationalising the full-stack export strategy through industry consortia and federal financing instruments, and was implemented within the deadline it set.

Exercised sovereignty is observable in a series of subsequent measures. A federal generative AI platform was made available to agencies from August 2025 and a defence-sector equivalent reported large-scale adoption within months. The Genesis Mission, established by executive order in November 2025, directs AI capabilities towards scientific discovery under the Department of Energy. A twenty-five per cent tariff on certain imported AI chips was adopted under Section 232 authority. The most consequential development concerns the vertical distribution of authority: the plan's recommendation to limit federal funding for states maintaining burdensome AI regulation was operationalised by an executive order of December 2025 establishing a litigation task force within the Department of Justice to challenge state AI statutes, instructing the Department of Commerce to identify excessively burdensome state laws and conditioning participation in a major broadband infrastructure programme. Legislative recommendations for federal preemption were transmitted to Congress in March 2026.

Assessments of the plan identify three limitations. It contains no federal commitment to protection against the harms that dominate the applied AI ethics literature, including fraud, discrimination, privacy violation and the generation of abusive material. Approximately one third of its actions do not identify a lead agency. And it generally lacks deadlines, funding pathways and resource allocations, which leaves its prospective effect uncertain despite the number of its recommendations. Santos Júnior and Brinkhues (2026), measuring the plan against practices identified in the peer-reviewed literature, find convergence in cybersecurity, technological sovereignty and standardisation alongside divergence in rights protection and environmental responsibility.

3.3 Cross-case comparison

Table 2 reports the comparison across the five dimensions, separating the declaration and instrumentation layers.

Table 2
Comparison of the Brazilian AI Plan and America's AI Action Plan Across Five Dimensions and Two Analytical Layers

Dimension	**PBIA 2024-2028: declared**	**PBIA 2024-2028: instrumented**	**AAP 2025: declared**	**AAP 2025: instrumented**
Technological autonomy	Technological and data sovereignty as premises; domestic AI processors; sovereign cloud	Half of structural investment in algorithms; 20% data; 20% compute; 1% cybersecurity; no structural allocation to connectivity	Unchallenged global dominance; leadership in the full stack	Domestic semiconductor manufacturing; export controls with location verification; high-security facilities

Dimension	PBIA 2024-2028: declared	PBIA 2024-2028: instrumented	AAP 2025: declared	AAP 2025: instrumented
Governance structure	Superior council, executive committee, five thematic chambers	Interministerial working group of fifteen bodies; monitoring externalised to an observatory; annual reporting	Federal coordination of more than ninety actions	Three concurrent executive orders; procurement conditionality; preemption of state regulation from December 2025
Ethical accountability	Algorithmic transparency centre; multisectoral ethics committee; regulatory framework by 2027	0.45% of the envelope; foreseen bodies not established; binding rules pending in the legislature since March 2025	Ideological neutrality and truth-seeking in federal AI	Revision of the NIST framework; procurement conditions; litigation against state statutes; no federal harm-protection commitment
Funding and infrastructure	R$ 23.03 billion for 2024-2028; supercomputer among the world's five most powerful	R$ 6.47 billion secured by December 2025, predominantly credit through Finep and BNDES; 2% of structural investment in electricity	Acceleration of infrastructure as a condition of leadership	Categorical environmental exclusions; FAST-41 expansion; federal lands; grid policy; Section 232 tariff on certain imported chips
International positioning	Multilateral cooperation; Global South interlocution; ethical principles	Participation in governance fora; no dedicated export or alliance instrument	International AI diplomacy and security as a pillar	American AI Exports Program; plurilateral alliances; conditioning of rivals' access to compute

Note. Percentages refer to the announced envelope in the case of the axis allocations and to structural investment in the case of the enabler allocations. Source: prepared by the author, based on the documents and secondary sources described in the method section.

Three patterns emerge from the comparison. The first concerns the direction of the translation gap. In the Brazilian case, declaration exceeds instrumentation in every dimension: sovereignty is claimed transversally

and funded selectively, with the governance dimension exhibiting the widest divergence between what is promised and what is allocated. In the American case the relation is inverted in one dimension and preserved in the others. Ethical accountability is instrumented beyond what is declared, since the plan states a narrow commitment to neutrality in federal systems and produces instruments that reach the regulatory autonomy of subnational units. The gap is not a uniform property of policy documents but varies in direction.

The second pattern concerns the layers of the stack that each plan addresses. Brazilian instrumentation concentrates on layers with low capital intensity and high domestic appropriability, namely algorithms, data, education and, in declaration, regulation. American instrumentation centres on layers with high capital intensity and contested international access, namely semiconductors, energy, permitting and compute distribution. Neither plan addresses the layers the other prioritises with comparable weight.

The third pattern concerns the direction in which the principal instrument operates. Brazilian instruments operate inward, seeking to form capability within the national territory. American instruments operate in both directions, but the pillar that carries the plan's distinctive content operates outward, conditioning what other states may acquire and on what terms. The Brazilian plan contains no instrument of this kind, and the absence is structural rather than a matter of choice, since conditioning the access of others requires possession of assets others need.

4. DISCUSSION

The two plans differ in the translation from claim to instrument in a manner that follows from their positions in the international division of AI production. The Brazilian plan claims sovereignty comprehensively and instruments it partially, concentrating resources on layers where domestic

capability can be formed at accessible cost and leaving the layers of highest dependence addressed by declaration alone. The American plan claims dominance and instruments it through the removal of domestic constraints and the conditioning of external access, leaving the accountability functions that the governance literature treats as constitutive of legitimate control unaddressed at federal level and contested at subnational level.

The divergence is not reducible to differences in resources. A plan with a larger envelope would not by itself give Brazil an outward instrument, because such instruments depend on possessing assets that other states require. Nor would a smaller envelope have prevented the United States from adopting an accountability agenda, since the instruments it deployed for other purposes were available for that one. What varies with position is the repertoire of instruments through which sovereignty can be pursued, and what varies with political choice is the selection made within that repertoire.

4.1 A typology of sovereignty strategies

The comparison supports the designation of two strategies, which the framework of the article allows to be situated within a broader space. Adaptive sovereignty treats autonomy as a capability under construction, pursued through developmental instruments, international cooperation and alignment with circulating ethical principles, on the premise that norm participation is a form of influence available to states that cannot compete on capability. Propositive sovereignty treats autonomy as a position to be defended and projected, pursued through the removal of domestic constraints, the differentiation of allies from rivals and the establishment of national configurations as international defaults.

The two strategies occupy adjacent cells of a space defined by two dimensions. The first is the direction of the principal instrument, which may operate inward, forming capability within the jurisdiction, or outward,

conditioning the access of others. The second is the locus of leverage, which may derive from possession of upstream assets or from jurisdictional and market authority. Figure 2 presents the resulting typology.

Figure 2
Typology of Digital Sovereignty Strategies by Direction of Instrument and Locus of Leverage

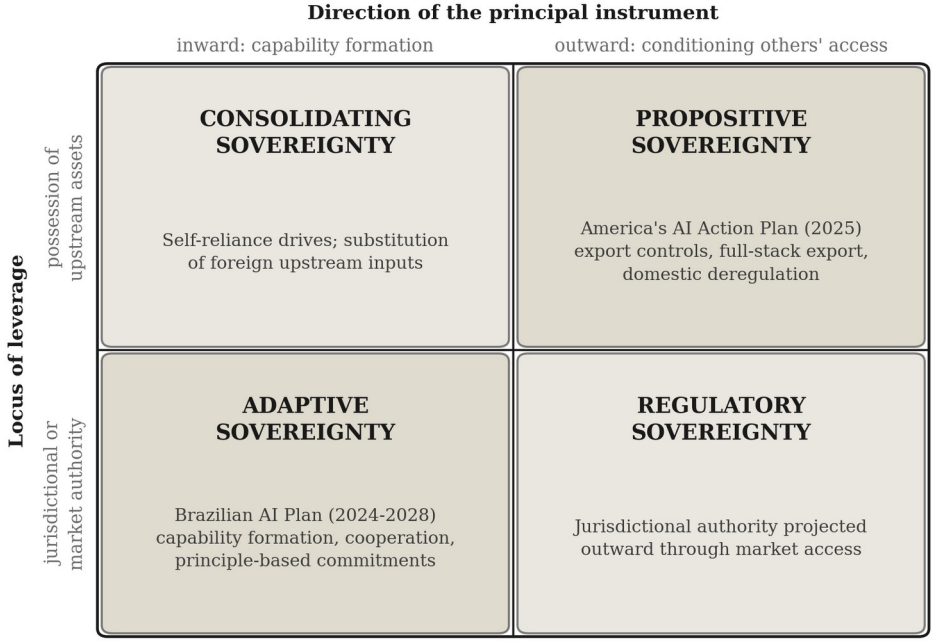

Note. The cells occupied by the two cases analysed are shaded. Source: prepared by the author, extending Papyshev and Yarime (2023).

Two further cells follow from the dimensions and situate the cases analysed. Consolidating sovereignty combines possession of upstream assets with inward instruments, and describes self-reliance drives oriented to the substitution of foreign inputs, a configuration documented for the Chinese case (Hung, 2025; Roberts et al., 2021) and characteristic of American policy before export restriction became a principal instrument. Regulatory sovereignty combines jurisdictional authority with outward projection, and describes the European configuration in which market access is the

mechanism through which domestic rules acquire extraterritorial effect (Bradford, 2012), a configuration whose limits Calderaro and Blumfelde (2022) and Mügge (2024) examine.

The typology extends the state-role triad of Papyshev and Yarime (2023) rather than replacing it. Their categories of development, control and promotion describe functions performed within the jurisdiction and therefore populate the inward column: development corresponds to adaptive and consolidating strategies, promotion to the domestic component of propositive strategies, and control to the domestic component of regulatory strategies. The outward column has no equivalent in the triad, because the corpus from which it was derived predates the period in which conditioning the access of other states became a principal instrument of AI policy. Adding the outward axis preserves the explanatory value of the original categories while accommodating instruments they were not constructed to observe.

4.2 Propositions

Three propositions follow from the comparison. They are stated as conditional explanations rather than as tested findings, and each carries the boundary conditions under which it is expected to hold.

The first concerns the alignment between structural position and the layers that instruments address. The greater a state's dependence on foreign upstream layers, the more its sovereignty instruments concentrate on layers with low capital intensity and high domestic appropriability, and the less on layers with high capital intensity and internationally contested access. The Brazilian distribution supports this reading: close to half of structural investment addresses algorithms, where domestic scientific capacity exists, while electricity receives two per cent and cybersecurity one per cent, and connectivity receives no dedicated structural allocation (Johansson Neto et al., 2024). The mechanism proposed is not fiscal alone. Investment in high-capital layers yields returns only if access to

complementary inputs is secure, and access to those inputs is precisely what the dependent position does not guarantee, so that concentration on appropriable layers is a rational adjustment to constraint rather than a planning deficiency. The proposition is expected to hold for states outside the technological core that retain a functioning research base and fiscal capacity for public investment. It is not expected to hold for states whose leverage derives from energy or land endowments, where investment in high-capital layers may be secured by the endowment itself.

The second concerns the distribution of accountability. Declared accountability commitment and institutional capacity to enforce it are inversely distributed across structural positions. Adopter states declare extensive accountability agendas they have not yet institutionalised; producer states possess the institutional capacity to institutionalise such agendas and redirect it away from them. The Brazilian case exhibits the first pattern, with an algorithmic transparency centre foreseen but not created, an allocation of 0.45 per cent to the corresponding axis and a binding framework pending in the legislature since March 2025. The American case exhibits the second, with the revision of the NIST framework, the establishment of ideological neutrality as a procurement criterion and the December 2025 initiative constraining the subnational statutes that had begun to supply accountability functions. The mechanism proposed is that the perceived cost of oversight varies with position: for a state competing at the frontier, oversight is experienced as loss of relative speed, whereas for a state building capability it is experienced as a claim to legitimacy that costs little at the point of declaration. This is a departure from the implementation gap identified by Corrêa et al. (2023) and Birkstedt et al. (2023), which describes principles that fail to become rules. Active redirection is a distinct phenomenon, since it operates through institutions rather than in their absence, and the two should not be assimilated in comparative assessment.

The third concerns the interaction between the two strategies. When the producer state's principal outward instrument is an integrated full-stack export, adoption by a client state resolves its capability constraint at the cost of the autonomy that inward capability formation seeks to build. Adaptive sovereignty and stack adoption are therefore substitutes rather than complements, and the window available to the adopter for capability formation narrows as stack integration deepens. The evidence is indirect and consistent rather than demonstrative: Executive Order 14320 packages hardware, models, software and standards for export as a unit, and the Brazilian sovereign cloud initiative contemplates contracting foreign providers for services the initiative was intended to internalise. The mechanism proposed is that integration across layers raises switching costs, so that adoption at one layer propagates dependence to adjacent layers, in line with the infrastructural dynamics documented by van der Vlist et al. (2024). The proposition is expected to hold where the exported stack is integrated across layers and where the adopting state lacks a domestic alternative at any layer. It is expected to weaken where open-weight models and modular procurement allow layer-by-layer substitution, which identifies a policy lever available to adopter states and returns theoretical interest to the digital public infrastructure pathway proposed by Sengupta et al. (2025).

Table 3 summarises the propositions, the evidence on which each rests and its boundary conditions.

Table 3
Propositions Derived From the Comparison, Supporting Evidence and Boundary Conditions

Proposition	Statement	Supporting evidence	Boundary conditions
P1. Instrument-position alignment	Dependence on foreign upstream layers shifts instruments towards layers of low capital intensity and high domestic appropriability	Brazilian structural allocation: algorithms ~50%, data ~20%, compute ~20%, education 11%, electricity 2%, cybersecurity 1%, connectivity none (Johansson Neto et al., 2024)	Holds for states outside the technological core with a functioning research base and fiscal capacity; weakens where leverage derives from energy or land endowments
P2. Inverse distribution of accountability	Declared accountability and capacity to enforce it are inversely distributed across structural positions	Brazil: 0.45% to the governance axis, transparency centre not created, binding framework pending since March 2025. United States: NIST framework revision, procurement neutrality conditions, preemption of state statutes from December 2025	Holds where the producer state faces frontier competition perceived as time-sensitive; weakens where domestic constitutional constraints protect subnational regulatory authority
P3. Substitution between adaptive sovereignty and stack adoption	Adoption of an integrated exported stack resolves capability constraints at the cost of the autonomy that inward capability formation seeks	Executive Order 14320 packages hardware, models, software and standards as a unit; the Brazilian sovereign cloud contemplates foreign providers for internalised services	Holds where the exported stack is integrated and no domestic alternative exists at any layer; weakens where open-weight models and modular procurement permit layer-by-layer substitution

Note. The propositions are conditional explanations derived retroductively; they are not tested findings. Source: prepared by the author.

4.3 Alternative explanations and relation to prior findings

Three alternative accounts of the observed divergence deserve consideration, since the design does not exclude them.

The first attributes the divergence to fiscal capacity alone. On this reading, Brazil would instrument sovereignty as the United States does if it

possessed comparable resources, and position adds nothing to the explanation. The account is incomplete for two reasons. Outward instruments require assets other states need, which resources alone do not create, and the accountability pattern documented in the American case runs opposite to what resource abundance would predict, since the capacity to institutionalise oversight was available and was directed elsewhere. Resources condition the scale of instrumentation, not its direction.

The second attributes the divergence to regulatory tradition, contrasting a statutory and precautionary European-influenced approach in Brazil with a market-oriented approach in the United States. This account has explanatory force for the governance dimension and is consistent with the state-role literature (Papyshev & Yarime, 2023; Radu, 2021). It does not account for the timing of the American shift, since the same tradition produced Executive Order 14110 in 2023 and its revocation in 2025, nor for the substitution dynamic described in the third proposition, which operates independently of the recipient state's regulatory tradition.

The third attributes the divergence to the electoral and ideological cycle, reading the American plan as the programme of a particular administration rather than as a structural response. The account is plausible and the plan's own genesis supports it, since it followed the revocation of the preceding framework. It is nonetheless insufficient on its own, because export restriction on advanced semiconductors preceded the change of administration and has been sustained across administrations, which indicates a structural component that the ideological account cannot absorb. The most defensible reading combines the two: the structural position defines the repertoire, and political orientation selects within it. The propositions above are formulated at the level of the repertoire, which is why they carry boundary conditions rather than predictions about particular governments.

The results converge with three established findings and qualify a fourth. They confirm that sovereignty aspiration frequently exceeds the industrial and institutional base required to sustain it, as Calderaro and Blumfelde (2022) establish for the European case and Ashraf and Veneziano (2026) generalise. They confirm the distance between declared principles and binding rules documented by Corrêa et al. (2023) and Birkstedt et al. (2023). They confirm the reading of state responses to AI concentration as an emergent international ordering rather than a defensive residue (Weymouth, 2025). And they are consistent with the finding of Schiavi et al. (2025) that Brazilian digitisation policy does not by itself sustain a sovereignty agenda, while adding that the shortfall is concentrated in identifiable enablers rather than distributed across the policy as a whole.

The qualification concerns the implementation gap. Comparative assessment has tended to treat the distance between principle and practice as a single phenomenon with a single cause, namely insufficient institutionalisation. The American case shows that an equivalent distance can arise from the opposite condition, where institutions exist and are redirected. Treating the two as instances of the same gap would obscure the difference between a state that has not yet built oversight capacity and a state that has built it and turned it towards other ends. The distinction matters for policy assessment, since the remedies differ, and for theory, since the mechanisms differ.

5. CONTRIBUTIONS

The theoretical contribution has three components, each supported by findings reported in Section 3.

The first is conceptual. The separation of declared, instrumented and exercised sovereignty supplies an intermediate category that the literature has lacked. Existing work has distinguished sovereignty as discourse from sovereignty as capability (Pohle & Thiel, 2020; Roberts, 2024), a distinction

that leaves the analyst with rhetoric on one side and outcomes on the other and no way of examining the selection of instruments that connects them. Making instrumentation a distinct layer allows the translation gap to be observed as a property of a policy configuration, with a direction and a magnitude that vary across dimensions. The Brazilian case shows a gap that is uniform in direction and widest in the governance dimension; the American case shows a gap that inverts in one dimension, where instrumentation exceeds declaration. Neither observation is available to a two-term distinction.

The second is typological. The comparison supports extending the state-role triad of Papyshev and Yarime (2023) by adding a dimension for the direction in which the principal instrument operates. The extension is warranted by evidence rather than by conceptual preference: the third pillar of the American plan and Executive Order 14320 instrument sovereignty by conditioning what other states may acquire, and this instrument cannot be classified as development, control or promotion without losing what distinguishes it. The resulting four-cell typology retains the original categories in the inward column and accommodates in the outward column both the American configuration and the European regulatory configuration that Bradford (2012) describes through a different mechanism. The extension follows the requirement that a theoretical contribution alter what is known rather than confirm it (Whetten, 1989), and its scope is delimited by the period from which the evidence is drawn.

The third is explanatory. The three propositions specify mechanisms that connect structural position to instrument selection, with boundary conditions that make them examinable in other cases. The first identifies capital intensity and appropriability as the properties that determine which layers a dependent state addresses. The second identifies the perceived cost of oversight as varying with competitive position, and distinguishes active redirection from the implementation gap. The third identifies integration

across layers as the mechanism through which stack adoption forecloses capability formation, and specifies modularity as the condition under which the foreclosure weakens.

The contribution is bounded in two respects. It rests on two cases and supports analytical rather than statistical generalisation. And it addresses the relation between position and instrument, not the effect of instrument on technological outcome, which the design cannot establish.

The coding protocol constitutes a modest methodological contribution. Comparative analysis of national AI strategies has relied predominantly on the textual content of the documents, whether through qualitative content analysis or topic modelling (Hung, 2025; Papyshev & Yarime, 2023; Saheb & Saheb, 2024). Textual methods observe declaration well and instrumentation poorly, since a document may devote extensive text to an objective it funds marginally. Combining textual indicators of claim with allocative and legal indicators of instrument, and assigning each segment to a layer by an explicit rule, makes the divergence between the two observable within the analysis rather than introduced by the analyst's judgement. The protocol is transferable to other national plans wherever allocative data or implementing acts are available, and its principal requirement, a secondary corpus recording obligations and allocations, is met by most contemporary strategies.

The findings carry implications for policy design that follow from the demonstrated results and should be distinguished from the contributions above, which concern knowledge rather than action.

For the Brazilian case, three points follow. The concentration of financing in credit instruments, which the execution data confirm as the dominant channel, addresses firms' cost of capital but not the demand-side uncertainty characteristic of frontier innovation, and public procurement for innovation would supply an instrument the plan currently lacks. The near-absence of provisions on energy and connectivity is consequential beyond

the sums involved, since the terms on which foreign data centre investment is admitted are negotiated at exactly this junction, and the country's position in renewable generation is an asset that the plan does not convert into leverage. And the institutions foreseen under the governance axis have a marginal cost relative to the envelope, so that their non-establishment weakens the plan's normative claims at little fiscal saving. The third proposition adds a design criterion that the plan does not currently apply: where procurement can be modularised and open-weight models substituted at specific layers, dependence propagates less readily across the stack.

For adopter states more generally, the analysis indicates that the layers over which leverage is retained are those where public authority is not readily substitutable by foreign provision, which include public service delivery infrastructure, public data assets, regulation and human capital. The digital public infrastructure pathway examined by Sengupta et al. (2025) is the clearest instance, and the results here give it a structural rationale rather than a normative one.

For assessment practice, whether by oversight bodies, funders or researchers, the layered framework supplies a criterion. A plan should be evaluated not by the ambition of its claims nor by outcomes it has not had time to produce, but by the alignment between what it claims and what it funds, an alignment that is observable early and that predicts little about intentions while revealing much about priorities.

FINAL CONSIDERATIONS

This article asked how the Brazilian AI Plan and America's AI Action Plan differ in translating digital sovereignty claims into policy instruments, and what that divergence reveals about the relation between structural position in AI production and the sovereignty strategy available to a state. The answer is that the two plans diverge systematically at the instrumentation layer, and that the divergence follows the repertoire of

instruments that each position makes available rather than the content of the claims each document makes.

The three specific objectives were met by identifiable blocks of results. The first, characterising declared sovereignty, is answered in Sections 3.1 and 3.2, which show a transversal claim in the Brazilian case and an expansive claim of dominance in the American case. The second, identifying instruments and locating the divergence, is answered in the same sections and in Table 2, which record an allocation of 0.45 per cent to the Brazilian governance axis against a transversal sovereignty claim, a concentration of structural investment in algorithms with two per cent in electricity, and, in the American case, an instrumentation of accountability that runs beyond the narrow commitment declared and reaches subnational regulatory authority. The third, deriving propositions, is answered in Section 4.2 and Table 3.

The contributions sustained by these results are the three-layer conceptual distinction, the extension of the state-role typology to accommodate sovereignty exercised over the access of others, the three propositions with their boundary conditions, and the coding protocol that renders the translation gap observable. Each rests on findings reported in Section 3, and none extends beyond what a two-case documentary design can support.

The conclusions apply under stated conditions. They concern national AI plans issued by executive authority in the period following the diffusion of generative models, in states occupying either a producer or an adopter position in the AI value chain. They do not extend to states whose leverage derives principally from energy or land endowments, nor to jurisdictions whose regulatory authority produces extraterritorial effects through market access, both of which the typology accommodates as distinct cells requiring separate examination.

Three limitations bound the inferences. Documentary analysis captures institutional design and allocative choice, not implementation dynamics, bureaucratic behaviour or the informal arrangements through which policy is enacted, which means that the exercised layer is assessed here only through reported execution data. The absence of interviews precludes access to the deliberative processes that produced the documents and to the interpretations guiding their execution. And both plans remain in course, so that the execution gap is measured at a single point and may narrow or widen.

Each limitation defines a research avenue. Longitudinal analysis of the indicators published by the Brazilian Artificial Intelligence Observatory would allow the translation and execution gaps to be tracked rather than observed once, and would test whether the Brazilian allocation pattern persists as resources are released. Interview-based research with actors in the plan's working group and in the parallel legislative process would reach the coordination mechanisms that documents do not record. Extension of the framework to additional cases, in particular to middle-income states pursuing autonomy under comparable constraints and to states whose leverage derives from energy endowments, would establish whether the adaptive configuration identified here is a general pattern or a national one, and would populate the cells of the typology that the present evidence leaves as inference. The third proposition invites a further inquiry with practical stakes: whether modular procurement and open-weight substitution measurably slow the propagation of dependence across stack layers, a question that observation of adopter states over the coming years will make answerable.

Digital sovereignty, under either model, is not achieved by declaration. It depends on the capacity to align infrastructure, financing and regulation and to sustain that alignment across political cycles. Measured against that criterion, the two plans examined here face different tests. Brazil must

convert an announced envelope into executed capability at the layers where dependence is greatest. The United States must show that arrangements resting on executive discretion can outlast the discretion that produced them.

REFERENCES

Ashraf, A., & Veneziano, V. (2026). AI sovereignty without power? Reviewing strategic agency, infrastructural dependence and digital imperialism. AI & Society. Advance online publication. https://doi.org/10.1007/s00146-026-03116-4

Bareis, J., & Katzenbach, C. (2022). Talking AI into being: The narratives and imaginaries of national AI strategies and their performative politics. Science, Technology, & Human Values, 47(5), 855-881. https://doi.org/10.1177/01622439211030007

Belli, L. (2023). Exploring the key AI sovereignty enablers (KASE) of Brazil, towards an AI sovereignty stack. SSRN. https://doi.org/10.2139/ssrn.4465501

Bhaskar, R. (2008). A realist theory of science. Routledge. https://doi.org/10.4324/9780203090732

Birkstedt, T., Minkkinen, M., Tandon, A., & Mäntymäki, M. (2023). AI governance: Themes, knowledge gaps and future agendas. Internet Research, 33(7), 133-167. https://doi.org/10.1108/INTR-01-2022-0042

Bowen, G. A. (2009). Document analysis as a qualitative research method. Qualitative Research Journal, 9(2), 27-40. https://doi.org/10.3316/QRJ0902027

Bradford, A. (2012). The Brussels effect. Northwestern University Law Review, 107(1), 1-68.

Busuioc, M. (2021). Accountable artificial intelligence: Holding algorithms to account. Public Administration Review, 81(5), 825-836. https://doi.org/10.1111/puar.13293

Calderaro, A., & Blumfelde, S. (2022). Artificial intelligence and EU security: The false promise of digital sovereignty. European Security, 31(3), 415-434. https://doi.org/10.1080/09662839.2022.2101885

Chandra, Y., & Feng, N. (2026). Algorithms for a new season? Mapping a decade of research on the artificial intelligence-driven digital transformation of public administration. Public Management Review, 28(3), 620-654. https://doi.org/10.1080/14719037.2025.2450680

Corrêa, N. K., Galvão, C., Santos, J. W., Del Pino, C., Pinto, E. P., Barbosa, C., Massmann, D., Mambrini, R., Galvão, L., Terem, E., & de Oliveira, N. (2023). Worldwide AI ethics: A review of 200 guidelines and recommendations for AI governance. Patterns, 4(10), Article 100857. https://doi.org/10.1016/j.patter.2023.100857

Couture, S., & Toupin, S. (2019). What does the notion of "sovereignty" mean when referring to the digital? New Media & Society, 21(10), 2305-2322. https://doi.org/10.1177/1461444819865984

Danermark, B., Ekström, M., & Karlsson, J. C. (2019). Explaining society: Critical realism in the social sciences (2nd ed.). Routledge.

Djeffal, C., Siewert, M. B., & Wurster, S. (2022). Role of the state and responsibility in governing artificial intelligence: A comparative analysis of AI strategies. Journal of European Public Policy, 29(11), 1799-1821. https://doi.org/10.1080/13501763.2022.2094987

Dunleavy, P., & Margetts, H. (2025). Data science, artificial intelligence and the third wave of digital era governance. Public Policy and Administration, 40(2), 185-214. https://doi.org/10.1177/09520767231198737

Eisenhardt, K. M. (1989). Building theories from case study research. Academy of Management Review, 14(4), 532-550. https://doi.org/10.2307/258557

Filgueiras, F. (2022). Artificial intelligence policy regimes: Comparing politics and policy to national strategies for artificial intelligence. Global Perspectives, 3(1), Article 32362. https://doi.org/10.1525/gp.2022.32362

Fletcher, A. J. (2017). Applying critical realism in qualitative research: Methodology meets method. International Journal of Social Research Methodology, 20(2), 181-194. https://doi.org/10.1080/13645579.2016.1144401

Floridi, L. (2020). The fight for digital sovereignty: What it is, and why it matters, especially for the EU. Philosophy & Technology, 33(3), 369-378. https://doi.org/10.1007/s13347-020-00423-6

Hine, E., & Floridi, L. (2024). Artificial intelligence with American values and Chinese characteristics: A comparative analysis of American and Chinese governmental AI policies. AI & Society, 39(1), 257-278. https://doi.org/10.1007/s00146-022-01499-8

Hung, H. T. (2025). Exploring China's cyber sovereignty concept and artificial intelligence governance model: A machine learning approach. Journal of Computational Social Science, 8(1), Article 24. https://doi.org/10.1007/s42001-024-00346-8

Ishkhanyan, A. (2025). The sovereignty-internationalism paradox in AI governance: Digital federalism and global algorithmic control. Discover Artificial Intelligence, 5(1), Article 123. https://doi.org/10.1007/s44163-025-00374-x

Johansson Neto, G. P., Farias da Costa, V. C., & Gaspar, W. B. (2024). Brazil's Artificial Intelligence Plan (PBIA) of 2024: Enabler of AI sovereignty? The African Journal of Information and Communication, 34, 1-15. https://doi.org/10.23962/ajic.i34.20424

Mügge, D. (2024). EU AI sovereignty: For whom, to what end, and to whose benefit? Journal of European Public Policy, 31(8), 2200-2225. https://doi.org/10.1080/13501763.2024.2318475

Papyshev, G., & Chan, K. J. D. (2026). AI regulatory strategies for digital sovereignty: The role of geopolitics and technological disparities. Electronic Markets, 36(1), Article 8. https://doi.org/10.1007/s12525-025-00870-z

Papyshev, G., & Yarime, M. (2023). The state's role in governing artificial intelligence: Development, control, and promotion through national strategies. Policy Design and Practice, 6(1), 79-102. https://doi.org/10.1080/25741292.2022.2162252

Pohle, J., & Thiel, T. (2020). Digital sovereignty. Internet Policy Review, 9(4). https://doi.org/10.14763/2020.4.1532

Radu, R. (2021). Steering the governance of artificial intelligence: National strategies in perspective. Policy and Society, 40(2), 178-193. https://doi.org/10.1080/14494035.2021.1929728

Ribeiro, M. M., & Segatto, C. I. (2025). Artificial intelligence in Brazilian public organizations: Heterogeneities and capacity in information technology. Revista de Administração Pública, 59(1), Article e2024-0066. https://doi.org/10.1590/0034-761220240066x

Roberts, H. (2024). Digital sovereignty and artificial intelligence: A normative approach. Ethics and Information Technology, 26(4), Article 70. https://doi.org/10.1007/s10676-024-09810-5

Roberts, H., Cowls, J., Morley, J., Taddeo, M., Wang, V., & Floridi, L. (2021). The Chinese approach to artificial intelligence: An analysis of policy, ethics, and regulation. AI & Society, 36(1), 59-77. https://doi.org/10.1007/s00146-020-00992-2

Saheb, T., & Saheb, T. (2024). Mapping ethical artificial intelligence policy landscape: A mixed method analysis. Science and Engineering Ethics, 30(2), Article 9. https://doi.org/10.1007/s11948-024-00472-6

Santos Júnior, D. de B., & Brinkhues, R. A. (2026). America's AI action plan and global governance: Convergences, divergences, and theoretical

implications. Revista de Administração Pública, 60(1), Article e2025-0584. https://doi.org/10.1590/0034-761220250584x

Sayer, A. (2000). Realism and social science. Sage.

Schiavi, I., Silveira, S. A. da, & Pasti, A. (2025). Where does digital sovereignty lie in building the Brazilian smart city? Digital Geography and Society, 9, Article 100137. https://doi.org/10.1016/j.diggeo.2025.100137

Sengupta, A., Costa Barbosa, A., & Samdub, M. T. (2025). Understanding interrelationships between AI and digital public infrastructure (DPI) in India and Brazil. The African Journal of Information and Communication, 35, 1-11. https://doi.org/10.23962/ajic.i35.20800

Ulnicane, I., Knight, W., Leach, T., Stahl, B. C., & Wanjiku, W.-G. (2021). Framing governance for a contested emerging technology: Insights from AI policy. Policy and Society, 40(2), 158-177. https://doi.org/10.1080/14494035.2020.1855800

van der Vlist, F., Helmond, A., & Ferrari, F. (2024). Big AI: Cloud infrastructure dependence and the industrialisation of artificial intelligence. Big Data & Society, 11(1). https://doi.org/10.1177/20539517241232630

Weymouth, S. (2025). Digital disintegration: Techno-blocs and strategic sovereignty in the AI era. International Organization, 79(S1), S57-S70. https://doi.org/10.1017/S0020818325101070

Whetten, D. A. (1989). What constitutes a theoretical contribution? Academy of Management Review, 14(4), 490-495. https://doi.org/10.2307/258554

Wynn, D., Jr., & Williams, C. K. (2012). Principles for conducting critical realist case study research in information systems. MIS Quarterly, 36(3), 787-809. https://doi.org/10.2307/41703481

RESEARCH DATA AVAILABILITY STATEMENT

All data supporting the results of this study are published in the article itself. The documents analysed are publicly available: the Brazilian Artificial Intelligence Plan 2024-2028 is available on the portal of the Brazilian Ministry of Science, Technology and Innovation, and America's AI Action Plan on the portal of the White House. The normative acts, resolutions and execution reports composing the secondary corpus are identified in the method section and are open to public consultation.

FUNDING

This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

AUTHOR CONTRIBUTIONS

Darci de Borba Santos Júnior: Conceptualisation; Data curation; Formal analysis; Investigation; Methodology; Project administration; Visualisation; Writing - original draft; Writing - review and editing.

STATEMENT ON THE USE OF ARTIFICIAL INTELLIGENCE

Generative artificial intelligence tools were used to support the verification of bibliographic metadata and language editing. The formulation of the research problem, the methodological design, the analysis of the documents, the derivation of the propositions and the final text are the sole responsibility of the author, who reviewed and validated all the content presented.

CONFLICT OF INTEREST STATEMENT

The author declares that there is no conflict of interest to report.

AUTHOR BIOGRAPHY

Darci de Borba Santos Júnior holds a doctorate and is a researcher at the Institute for Applied Economic Research (Ipea), affiliated with the Graduate Programme in Public Policy and Development. His work addresses digital transformation, public management, artificial intelligence, governance and process management, following a career in financial institutions, public administration, consultancy and teaching.

This preprint was submitted under the following conditions:

- The authors declare that the necessary Terms of Free and Informed Consent of participants or patients in the research were obtained and are described in the manuscript, when applicable.
- The authors declare that the preparation of the manuscript followed the ethical norms of scientific communication.
- The authors declare that they are aware that they are solely responsible for the content of the preprint and that the deposit in SciELO Preprints does not mean any commitment on the part of SciELO, except its preservation and dissemination.
- The authors declare that the data, applications, and other content underlying the manuscript are referenced.
- The deposited manuscript is in PDF format.
- The authors declare that the research that originated the manuscript followed good ethical practices and that the necessary approvals from research ethics committees, when applicable, are described in the manuscript.
- The authors declare that once a manuscript is posted on the SciELO Preprints server, it can only be taken down on request to the SciELO Preprints server Editorial Secretariat, who will post a retraction notice in its place.
- The authors agree that the approved manuscript will be made available under a Creative Commons CC-BY license.
- The submitting author declares that the contributions of all authors and conflict of interest statement are included explicitly and in specific sections of the manuscript.
- The authors declare that the manuscript was not deposited and/or previously made available on another preprint server or published by a journal.
- If the manuscript is being reviewed or being prepared for publishing but not yet published by a journal, the authors declare that they have received authorization from the journal to make this deposit.
- The submitting author declares that all authors of the manuscript agree with the submission to SciELO Preprints.