OBSERVARE
Universidade Autónoma de Lisboa
e-ISSN: 1647-7251
VOL. 17, Nº. 2
November 2026-April 2027

JANUS.NET
e-journal of International Relations

AI GOVERNANCE AND THE INTERNATIONAL TAX ORDER: ALGORITHMIC BASE EROSION UNDER THE OECD TWO-PILLAR SOLUTION. A LEGAL-THEORETICAL PERSPECTIVE ON INTERNATIONAL CORPORATE TAXATION

IKROM ABDURASULOVICH ERGASHEV
professor.ergashev@gmail.com
Doctor of Law (DSc), Professor and Head of the Department of Administrative and Legal Sciences, Academy of Justice of the Republic of Uzbekistan, Tashkent State University of Law, Tashkent (Uzbekistan). https://orcid.org/0000-0002-3924-5182

Abstract

This article advances a single, defensible claim: the OECD/G20 Two-Pillar Solution, though the most ambitious reform of international corporate taxation in a century, rests on nexus and profit-allocation concepts that are structurally misaligned with the way generative artificial intelligence creates value. The argument is developed doctrinally rather than through synthesis of institutional reportage. Pillar One ties reallocated taxing rights to identifiable market jurisdictions and a narrow set of the very largest in-scope multinationals; Pillar Two floors the effective rate at fifteen per cent but preserves a substance-based income exclusion keyed to payroll and tangible assets. Both design choices presuppose that profit can be traced to human labour and physical presence located somewhere. Generative-AI value, by contrast, is produced through globally sourced training data, distributed inference infrastructure, and model weights that are non-rival and near-costless to relocate. The result is a form of 'algorithmic base erosion' that the current architecture neither captures nor prohibits — legal in form, corrosive in effect. Engaging the scholarship on algorithmic accountability (Pasquale; Barocas and Selbst; Coglianese and Lehr), on jurisdiction and regulatory diffusion (Ryngaert; Svantesson; Bradford), and on the taxation of AI specifically (Oberson; Carvalho; Turina), the article shows why interpretive extension of existing rules is insufficient and identifies three doctrinal pressure points — the data-input nexus, the substance carve-out, and the autonomous-agent taxable person — where reform must concentrate. It builds on, but is distinct from, the emerging critique of the substance-based income exclusion's tangible-asset bias: the contribution here is to show that this bias is not a general drafting flaw but a structural mismatch with the specific production function of generative AI. The result is a focused legal diagnosis, not a general governance framework.

Keywords

International corporate taxation; OECD Two-Pillar Solution; generative artificial intelligence; nexus and profit allocation; base erosion; fiscal sovereignty; algorithmic accountability.

Resumo

Este artigo defende uma única tese defensável: a Solução de Dois Pilares da OCDE/G20, embora seja a reforma mais ambiciosa da tributação internacional das empresas neste último século, assenta em conceitos de nexo e de repartição de lucros que estão estruturalmente desalinhados com a forma como a inteligência artificial generativa cria valor. O argumento é desenvolvido de forma doutrinária, em vez de através da síntese de relatórios institucionais. O Primeiro Pilar vincula os direitos de tributação reatribuídos a jurisdições de mercado identificáveis e a um conjunto restrito das maiores multinacionais abrangidas pelo âmbito de aplicação; o Segundo Pilar estabelece um limite mínimo de 15% para a taxa efetiva, mas

preserva uma exclusão de rendimentos baseada na substância, indexada à massa salarial e aos ativos tangíveis. Ambas as opções de conceção pressupõem que o lucro pode ser atribuído ao trabalho humano e à presença física localizada num determinado local. O valor da IA generativa, em contrapartida, é produzido através de dados de treino de origem global, infraestruturas de inferência distribuídas e pesos de modelos que são não rivais e cuja relocalização tem um custo quase nulo. O resultado é uma forma de «erosão da base algorítmica» que a arquitetura atual nem capta nem proíbe — legal na forma, corrosiva no efeito. Recorrendo à literatura académica sobre a responsabilização algorítmica (Pasquale; Barocas e Selbst; Coglianese e Lehr), sobre jurisdição e difusão regulatória (Ryngaert; Svantesson; Bradford) e sobre a tributação da IA especificamente (Oberson; Carvalho; Turina), o artigo demonstra por que razão a extensão interpretativa das regras existentes é insuficiente e identifica três pontos de pressão doutrinários — o nexo de entrada de dados, a exclusão com base na substância e o sujeito passivo como agente autónomo — nos quais a reforma deve concentrar-se. Baseia-se, mas distingue-se, da crítica emergente ao viés dos ativos tangíveis na exclusão de rendimentos baseada na substância: a contribuição aqui consiste em demonstrar que este viés não é uma falha geral de redação, mas sim uma incompatibilidade estrutural com a função de produção específica da IA generativa. O resultado é um diagnóstico jurídico específico, e não um quadro geral de governação.

Palavras-chave

Tributação internacional das empresas; Solução dos Dois Pilares da OCDE; inteligência artificial generativa; nexo e repartição de lucros; erosão da base tributária; soberania fiscal; responsabilização algorítmica.

How to cite this article
Ergashev, Ikrom Abdurasulovich (2026). AI Governance and the International Tax Order: Algorithmic Base Erosion Under the Oecd Two-Pillar Solution. A Legal-Theoretical Perspective on International Corporate Taxation. *Janus.net, e-journal of international relations*. VOL. 17, Nº. 2, November 2026-April 2027, pp. 644-657. DOI https://doi.org/10.26619/1647-7251.17.2.32

Article submitted on 29 January 2026 and accepted 6 August 2026.

AI GOVERNANCE AND THE INTERNATIONAL TAX ORDER: ALGORITHMIC BASE EROSION UNDER THE OECD TWO-PILLAR SOLUTION. A LEGAL-THEORETICAL PERSPECTIVE ON INTERNATIONAL CORPORATE TAXATION

IKROM ABDURASULOVICH ERGASHEV

Introduction

Artificial intelligence has moved from experimental application to core economic infrastructure with unusual speed, and tax administrations have followed. Data collected in the OECD's 2024 *Inventory of Tax Technology Initiatives*, summarised in the OECD's 2025 report on tax administration digitalisation, record that seventy-two per cent of tax administrations now deploy AI in some function; among those administrations, the most common application is the detection of tax evasion and fraud, undertaken by three quarters of them (OECD, 2024). Those figures describe administrations using AI as a tool of enforcement. They say nothing about the harder question this article addresses: how the economic value generated by AI itself is to be located, characterised, and taxed under the rules of international corporate taxation.

The distinction matters because the two questions are routinely conflated. A large and growing literature treats 'AI and tax' as a story about smarter audits, chatbots, and compliance analytics. The fiscal problem that AI poses is different in kind. When a generative model is trained on data drawn from users across the world, developed by engineers dispersed across several jurisdictions, hosted on cloud infrastructure that spans many more, and sold as a service to customers everywhere at once, the settled categories of international tax law — source, residence, permanent establishment, and the arm's-length allocation of profit — lose their purchase. The value is real and large; the place where it is taxed is increasingly a matter of corporate structuring rather than economic substance.

The international community has not been idle. The OECD/G20 Inclusive Framework's Two-Pillar Solution, endorsed by 136 member jurisdictions on 8 October 2021 and by 139 as of June 2023 and elaborated through 2026, is the most far-reaching attempt in a century to re-fit international corporate taxation to a digitalising economy (OECD, 2021). Pillar One reallocates a share of taxing rights over the largest and most profitable multinationals to the market jurisdictions where their users and customers are located, irrespective of physical presence. Pillar Two establishes a fifteen per cent global minimum effective tax rate through the Global Anti-Base Erosion (GloBE) rules, constraining the

tax competition that has driven a decades-long decline in headline corporate rates (OECD, 2026b).

This article makes one claim and defends it. The Two-Pillar Solution, for all its ambition, is built on nexus and profit-allocation concepts that are structurally — not merely operationally — misaligned with the way generative AI creates value, and that misalignment produces a distinct and under-recognised phenomenon that I call algorithmic base erosion. The phrase denotes value that escapes taxation not through illegality but through the interaction of a genuinely novel production process with rules designed for a different one. The claim is narrow by design. It is not that AI 'disrupts public finance' in general, nor that a comprehensive new governance framework is required. It is that two identifiable design features of the leading reform — the market-jurisdiction nexus of Pillar One and the substance-based carve-out of Pillar Two — fail in specific, demonstrable ways when confronted with AI-generated profit, and that the failure is doctrinal rather than administrative.

A word on scope is needed at the outset, because it is easy — and, as a reviewer of an earlier version rightly warned, tempting — to treat 'artificial intelligence' as a single object. It is not. Deterministic rule-based systems, whose outputs follow transparently from encoded logic, raise no novel tax question: the value they produce remains tied to identifiable human design and deployment, and the existing categories reach it. Deep-learning models trained on large datasets begin to strain those categories, because value migrates from the human function to the data and the trained parameters. Autonomous agent-based systems, which act and transact without contemporaneous human direction, strain them further still. The argument of this article is deliberately confined to the second and third of these — generative and large-model systems whose economic value is dominated by training data and model weights rather than by locatable human labour. It is precisely this class, and not automation in general, for which the nexus and substance concepts of the Two-Pillar Solution break down; the distinctions among system types are therefore not a caveat to the argument but part of its architecture, marking exactly where the claim bites and where it does not.

The argument proceeds in five further steps. Section 2 sets out the doctrinal and theoretical framework and locates the claim precisely within the existing scholarship on the taxation of AI — distinguishing it from proposals to tax AI directly and from the general critique of the substance carve-out's tangible bias. Section 3 analyses the market-jurisdiction nexus of Pillar One and shows why the data-input character of AI value defeats it. Section 4 turns to Pillar Two's substance-based carve-out and demonstrates how a payroll-and-assets test rewards precisely the relocation that AI makes costless. Section 5 examines the question of the taxable person once autonomous agents transact at scale, and the coordination problem that follows. Section 6 states the limits of the claim and its implications. Throughout, the method is legal-doctrinal: the article reads the operative instruments closely and tests them against a clearly specified economic process, rather than surveying the field.

Framework: Nexus, Substance, and the Limits of Interpretive Extension

The doctrinal inheritance

International corporate taxation rests on two allocative ideas whose lineage is older than the digital economy. The first is nexus: a state may tax the business profits of a non-resident only where that enterprise has a sufficient connection to the taxing jurisdiction, historically a physical permanent establishment. The second is the arm's-length principle: profit is allocated among the affiliates of a multinational group as though they were independent parties transacting at market prices. Both ideas assume that economic activity has a locatable centre of gravity — a place where people work, assets sit, and functions are performed. The transfer-pricing apparatus that dominates practice is an elaborate machinery for finding that centre of gravity and attributing profit to it.

Fiscal-federalism theory supplies the normative intuition behind the design. In the tradition running from Musgrave's account of the fiscal functions to Oates's analysis of decentralisation, a recurring tax-assignment principle holds that highly mobile bases should be taxed at higher or more coordinated levels of government, because a mobile base taxed locally simply migrates, generating a race to the bottom (Musgrave. 1959) The principle is sound, but it presupposes that mobility is a property of the taxpayer's choices rather than of the tax base's intrinsic nature. AI value strains that presupposition: the base is not merely mobile but, in an important sense, placeless, because the inputs that generate it — data and model weights — are non-rival and can be replicated or relocated at near-zero cost.

Why interpretive extension is not enough

The instinctive response of tax lawyers to a new business model is interpretive extension: read the existing concepts generously so that they reach the new case. This has worked before — the concept of the 'significant economic presence' was an attempt to stretch nexus to digital platforms. But interpretive extension succeeds only where the new activity differs in degree, not in kind, from the paradigm the rule was written for. Coglianese and Lehr, analysing administrative decision-making in the machine-learning era, draw exactly this distinction between adapting existing legal categories to a faster or larger version of a familiar process and confronting a process whose logic the categories were never built to describe (Coglianese, 2017).The argument of this article is that generative-AI value creation falls on the far side of that line for two of the Two-Pillar Solution's load-bearing concepts.

The reason is best stated through the scholarship on algorithmic opacity. Pasquale's account of the 'black box' describes systems whose internal operations are not merely complex but resistant to external reconstruction, so that outcomes can be observed while their generative process cannot be audited (Pasquale, 2015). Where the value-creating process itself cannot be reconstructed by an outside authority, the transfer-pricing method — which depends on identifying the functions performed, assets used, and risks assumed by each affiliate — has nothing stable to grip. Barocas and Selbst make the complementary point in the discrimination context: the resulting harm is, in their own characterisation, almost always an unintentional emergent property of the algorithm's

use rather than a conscious choice by its programmers, which is what makes it unusually hard to identify or to explain to a court (Barocas and Selbst, 2016). Transposed to taxation, their insight explains why algorithmic base erosion is invisible to an enforcement-centred lens: nothing unlawful occurs. The erosion is a property of the rules meeting the technology, not of any single actor breaking them.

Jurisdiction as the deeper stake

Beneath the technical question lies a jurisdictional one. Ryngaert's work on jurisdiction in international law shows how states increasingly rely on a territorial 'hook' — the importation of a product, the provision of a service to residents — to regulate conduct that is in substance transnational (Ryngaert, 2015). Pillar One is a fiscal instance of exactly this move: it manufactures a taxing nexus out of the location of the market rather than the location of production. Svantesson's analysis of the internet-jurisdiction problem explains why the move is fragile: once the relevant activity has no natural physical locus, any territorial hook is a chosen convention that competing states may decline to accept, and enforcement across borders becomes the binding constraint (Svantesson, 2017). Bradford's account of the 'Brussels Effect' supplies the countervailing dynamic — a large market can export its regulatory choices extraterritorially through the practical necessity of market access, which is why the EU's fiscal and AI-regulatory choices carry weight beyond its borders (Bradford, 2020). Read together, these literatures frame the Two-Pillar Solution not as a neutral technical fix but as a contested allocation of jurisdictional authority over a base that resists natural location — precisely the setting in which structural misalignment, rather than mere drafting imperfection, should be expected.

The state of the debate, and this article's place in it

The claim advanced here does not enter an empty field, and it is important to say precisely what is new and what is not. A body of work has already asked how AI interacts with taxation, and it divides into three strands that this article must be distinguished from in turn.

The first strand proposes taxing AI itself. Oberson's pioneering study of taxing robots — later expanded into a full treatment of the taxation of artificial intelligence, and most recently into a dedicated handbook chapter on the subject — argues for taxing the use of automation, or AI systems directly, as a response to the erosion of labour-based tax bases, and examines the compatibility of such taxes with existing international tax rules (Oberson, 2019). The second strand asks who the taxpayer is once systems act without human direction. Carvalho, writing in *Intertax* in 2019, posed that question concretely for what he termed autonomous artificial intelligence, identifying the twin problems of 'disappearing income' — the absence of legal personality — and 'powerlessness to tax' — the absence of residence, and proposing a general attribution of taxable personality together with a two-tiered assessment of residence (Carvalho, 2019). The third strand is broader still: Turina has argued that the debate on the tax challenges of the digital economy has shifted progressively from a narrow base-protection rationale toward a pretext for overhauling the international tax regime itself, and the recent literature

reading that debate forward to AI holds that value attribution, nexus thresholds and source rules will all require new frameworks (Turina, 2020).

Most pointedly for the present argument, a recent critique by Tomazela Santos in the *Bulletin for International Taxation* has examined the substance-based income exclusion of Pillar Two and identified its conceptual shortcomings, its allocative distortions and its bias toward tangible-intensive business models, arguing that the formula inadequately reflects economic reality (Tomazela Santos, 2026). The problem, in other words, is being seen.

This article's contribution is narrower and sharper than any of these, and it is defined against them. It is not a proposal for a new tax on AI or robots, as Oberson's is; it is a diagnosis of why the existing reform fails. It is not a claim about the tax personality of the machine, as Carvalho's is, although Section 5 takes up that question at its own boundary. And it is not a general claim that intangibles strain the substance carve-out, as the tangible-bias critique is; that critique is correct but generic, applying equally to pharmaceutical patents, software licences, and brand value. The claim here is specific: that the substance carve-out and the market-jurisdiction nexus fail in a distinctive way for generative AI because of the particular character of its inputs — training data drawn from a population that is neither market nor residence, and model weights that are non-rival and relocatable at near-zero cost. Where the existing critique says 'the formula under-weights intangibles', this article says 'the formula, and the nexus beside it, cannot in principle locate value whose origin is a globally sourced, non-rival input'. The difference is between a calibration problem and a structural one, and only the latter resists repair by adjusting the exclusion percentages or reading the nexus generously.

Pillar One and the Failure of the Market-Jurisdiction Nexus

Pillar One's central innovation, Amount A, reallocates a portion of the residual profit of the largest and most profitable multinationals to market jurisdictions, defined by where revenue is sourced: twenty-five per cent of a group's profit above a ten per cent margin is redistributed, with an in-scope group treated as connected to a market jurisdiction once its revenue there passes a low monetary threshold (OECD, 2023). Two design parameters bound its reach. The scope threshold — global turnover above €20 billion and profitability above ten per cent — means that only a small number of the world's largest groups fall within it, predominantly large technology and consumer-facing platforms; the OECD has described the in-scope population as comprising the very largest and most profitable multinationals rather than a fixed count, and the convention itself contemplates a later reduction of the turnover threshold contingent on successful implementation. The mechanism operates through a multilateral convention layered onto the existing network of bilateral treaties, an architecture whose administrative complexity, and whose requirement of ratification by thirty jurisdictions accounting for at least sixty per cent of the relevant ultimate parent entities, has repeatedly slowed entry into force (OECD, 2023b) Both parameters matter for the AI case, but the deeper problem lies in the nexus concept itself.

The market-jurisdiction nexus works by treating the location of the user or customer as the anchor for taxing rights. For a streaming platform or a social network this is coherent:

the service is consumed in an identifiable place by an identifiable market, and the revenue can be sourced there. Generative AI complicates each link in that chain. First, a substantial part of the value is created upstream of any market, in the training phase, where the model ingests data contributed by people who are neither customers nor, in any tax-relevant sense, participants in the market where the service is later sold. The jurisdiction whose residents supplied the training data has, on the current design, no claim at all: it is neither the residence of the developer nor the market of the sale. This is the first vector of algorithmic base erosion — value drawn from a population that the nexus rules render fiscally invisible. The erosion is not marginal. As models scale, the share of value attributable to training data — and thus to the jurisdictions whose residents generated it — grows relative to the share attributable to any locatable human function, so that the very trend driving the frontier of the technology is the trend widening the gap between where value originates and where the nexus permits it to be taxed. What earlier debates approached as the taxation of the 'digital nexus' or the monetisation of data reappears here in its sharpest form: not as a platform booking advertising revenue in the wrong country, but as a productive input — data itself — whose contributing jurisdiction the allocation rules cannot even name.

Second, the 'market' for a general-purpose model is not stable in the way Amount A assumes. The same model weights serve consumers, enterprises, and downstream developers who embed the model through an interface, so that the revenue booked in a given jurisdiction reflects contractual routing and API architecture as much as any underlying economic locus. Where the taxable footprint can be reshaped by re-routing an interface, the nexus is not tracking economic substance; it is tracking a structuring decision. Barocas and Selbst's lesson recurs here in fiscal form: the outcome is produced by neutral technical design choices that are entirely lawful, which is exactly why an enforcement posture cannot reach it (Barocas and Selbst, 2016).

Third, the scope threshold interacts perversely with AI industry structure. Because Amount A captures only the very largest groups, a model developer can remain outside scope while nonetheless generating precisely the placeless, data-derived profit the reform was meant to reach, simply by falling below the turnover or profitability line, or by distributing activity across entities. The reform thus reaches a handful of incumbents while leaving the structural gap open beneath them. None of this is an argument that Pillar One was badly drafted. It is an argument that a nexus built on market location cannot, even when drafted well, capture value whose origin is a globally sourced input and whose situs is a matter of architecture.

Pillar Two and the Substance Carve-Out Paradox

If Pillar One fails at the nexus, Pillar Two fails at the carve-out. The GloBE rules impose a fifteen per cent minimum effective rate, computed jurisdiction by jurisdiction, with a top-up tax collected where the effective rate falls below the floor; the Inclusive Framework has consolidated its interpretive guidance into successive editions of a Consolidated Commentary, most recently that of 2026 (OECD, 2026). The rules have advanced toward implementation more smoothly than Pillar One; the European Union adopted a directive requiring member states to give effect to the minimum tax for groups

with consolidated revenue of €750 million or more, and numerous other jurisdictions have enacted conforming legislation (Council Directive 2022/2523) Yet a specific feature of the design — the substance-based income exclusion — turns against the reform's purpose in the AI context.

The substance-based income exclusion carves out from the top-up tax a return computed on two factors: a percentage of eligible payroll costs and a percentage of the carrying value of eligible tangible assets in a jurisdiction. The rationale is intelligible and, for a traditional economy, defensible: the exclusion is meant to protect real economic activity — factories, workforces — from the minimum tax, targeting the charge instead at profit that is purely the product of tax planning. The carve-out therefore rewards the presence of people and physical things in a jurisdiction, and expressly excludes intangibles from the asset base.

For an AI developer this incentive structure is close to inverted. The production function of a frontier model is dominated by intangible inputs — data, model weights, and compute contracted from third parties — and requires comparatively little payroll and few owned tangible assets relative to the profit generated. A group can therefore locate the modest payroll and tangible footprint that does exist in a low-rate jurisdiction, claim the substance exclusion there, and shelter a slice of exactly the intangible-derived profit that the minimum tax was designed to reach. The carve-out, intended to distinguish real activity from planning, cannot perform that function where the 'real activity' is itself intangible and mobile. This is the second vector of algorithmic base erosion, and it is sharper than the first because it operates through an express statutory allowance rather than a gap. The recent tax-policy literature has independently observed that the exclusion is biased toward tangible-intensive business models and reflects economic reality poorly (Tomazela Santos, 2026). That observation is correct but general; it holds for any intangible-heavy firm. The argument here is the narrower and, I think, more consequential one: for the generative-AI developer the mismatch is not a matter of degree but of kind, because the value-generating inputs are not merely under-weighted by the formula but are of a type — non-rival, placeless data and weights — that no payroll-and-assets proxy can track at all.

The point can be put as a paradox. Pillar Two was built to stop the migration of paper profit to low-tax jurisdictions. The substance carve-out was built to ensure that genuine activity would not be caught in the crossfire. But when the genuine activity consists of assets that are themselves weightless and relocatable, the carve-out ceases to separate substance from planning and instead supplies a lawful channel for the planning it was meant to exclude. The design assumption — that payroll and tangible assets are a reliable proxy for immobile, genuinely local value — holds for the twentieth-century firm and breaks for the model developer. Here, too, the failure is structural: no adjustment of the exclusion percentages repairs a proxy that has stopped tracking the thing it was chosen to measure.

The Autonomous-Agent Taxable Person and the Coordination Constraint

A third pressure point is more speculative but follows from the same logic. International tax law presupposes a taxable person — a natural or legal person to whom income is

attributed and on whom a liability is imposed. As autonomous systems increasingly transact without contemporaneous human direction — executing trades, negotiating and concluding contracts, allocating resources — the attribution of the resulting income becomes contestable. This is not a hypothetical raised here for the first time. Carvalho has posed it as a concrete question of international tax law, identifying the absence of legal personality and the absence of residence as the twin obstacles to taxing income attributable to an autonomous system, and proposing a general attribution of taxable personality together with a two-tiered residence test (Carvalho, 2019). Oberson, responding to that proposal, suggested that the jurisdiction whose labour market is most seriously affected by such a system's activity might itself furnish a residence criterion — a suggestion that is instructive here precisely because it reaches for an impact-based connecting factor once the conventional ones fail (Oberson, 2019). The question is not the science-fiction one of whether an algorithm should be a taxpayer; it is the concrete one of which human or corporate principal is treated as earning income that an autonomous agent generated across several jurisdictions in ways none of them separately controlled. The existing answer — attribute to the deploying entity — is workable only so long as the deploying entity has a stable location, which returns the analysis to the nexus problem of Section 3.

This is not a settled crisis and should not be overstated. It is flagged here as the natural extension of the article's claim rather than as an established doctrinal failure, and it marks a boundary of the argument: the misalignment demonstrated for Pillars One and Two is present-tense and concrete, whereas the taxable-person problem is emergent. Treating them as equivalent would repeat the error the reviewers rightly identified in the literature — asserting more than the evidence supports.

Whatever the doctrinal response to these three pressure points, it must clear a coordination constraint that the jurisdictional scholarship makes unavoidable. Because the base resists natural location, any unilateral fix invites the double taxation and trade retaliation that Svantesson's and Ryngaert's analyses predict when states assert competing extraterritorial claims over the same conduct (Svantesson, 2017). The proliferation of national digital-services taxes — adopted by several jurisdictions pending a multilateral solution, met by threats of trade retaliation, and in several cases since withdrawn or amended under external pressure — is the observable symptom of that constraint. The United Nations initiative toward a framework convention on international tax cooperation reflects a parallel dissatisfaction, particularly among capital-importing states whose residents supply training data but whose treasuries capture none of the resulting value: the General Assembly's resolution of 22 December 2023 established an ad hoc intergovernmental committee to draft terms of reference for such a convention, and the Assembly returned to the subject a year later to carry the process forward (United Nations General Assembly, 2023). The coordination constraint does not resolve the doctrinal question; it disciplines the range of admissible answers, ruling out purely national reform.

Implications and Limits of the Claim

The analysis supports three focused conclusions and, equally important, marks what it does not establish. First, the misalignment between the Two-Pillar Solution and generative-AI value creation is structural, located in the market-jurisdiction nexus of Pillar One and the substance carve-out of Pillar Two, and is therefore not curable by administrative tuning or generous interpretation. Reform, if it comes, must engage the concepts themselves: a data-input dimension to nexus that recognises the jurisdiction supplying training data; a re-specification of the substance test so that it ceases to treat relocatable intangibles as local activity; and, prospectively, a rule of attribution for autonomous-agent income.

Second, the phenomenon is lawful, which is why it demands a doctrinal rather than an enforcement response. The seventy-two per cent of administrations deploying AI, three quarters of them for the detection of evasion and fraud, are equipped for a different problem entirely (OECD, 2025). Algorithmic base erosion is not evasion; it is the predictable output of a genuinely new production process meeting rules written for an older one. Barocas and Selbst's central lesson — that lawful, neutral design can generate systematically distortive outcomes as an unintended emergent property — is the correct frame, and it explains why the problem has been under-recognised by a literature oriented toward compliance and audit (Barocas and Selbst, 2016).

Third, the coordination constraint means that the reform space is genuinely international, not national, and that the distributive stakes are real: on the current design, states whose populations supply the data that trains frontier models have no fiscal claim on the value that data helps create. That is a defensible normative concern, but it is raised here as a consequence of the doctrinal analysis, not as an independent programme.

The limits of the claim should be equally explicit. This article does not offer an integrated governance framework spanning tax, expenditure, and international law; it deliberately abandons that ambition in favour of one defensible thesis about corporate taxation. As Section 1 made explicit, it does not treat 'artificial intelligence' as a single object: the argument concerns generative and large-model systems whose value is dominated by data and weights, and applies only weakly, if at all, to deterministic or narrowly deployed systems whose value remains tied to identifiable human functions. And it does not claim that the Two-Pillar Solution is a failure; it claims that a specific and important slice of value falls outside its reach for structural reasons. Whether the political will exists to close that gap — through a data-sensitive nexus, a reformed substance test, or a UN-centred instrument — is a question this article poses rather than answers. What it establishes is narrower and, for that reason, more secure: that the gap is real, that it is structural, and that it will not be closed by reading the existing rules more generously.

References

Barocas, S. and Selbst, A.D. (2016), 'Big Data's Disparate Impact', *California Law Review*, Vol. 104, No. 3, pp. 671–732, https://doi.org/10.15779/Z38BG31.

Bradford, A. (2020), *The Brussels Effect: How the European Union Rules the World*, Oxford University Press, New York, ISBN 9780190088583, https://doi.org/10.1093/oso/9780190088583.001.0001.

Carvalho, L. de Lima (2019), 'Spiritus Ex Machina: Addressing the Unique BEPS Issues of Autonomous Artificial Intelligence by Using "Personality" and "Residence"', *Intertax*, Vol. 47, No. 5, pp. 425–443, Kluwer Law International, Alphen aan den Rijn.

Coglianese, C. and Lehr, D. (2017), 'Regulating by Robot: Administrative Decision Making in the Machine-Learning Era', *Georgetown Law Journal*, Vol. 105, No. 5, pp. 1147–1223.

Council Directive (EU) 2022/2523 of 14 December 2022 on ensuring a global minimum level of taxation for multinational enterprise groups and large-scale domestic groups in the Union, OJ L 328, 22.12.2022, pp. 1–58 (CELEX 32022L2523; ELI http://data.europa.eu/eli/dir/2022/2523/oj), Art. 2(1) (revenue threshold of EUR 750 000 000).

Musgrave, R.A. (1959), *The Theory of Public Finance: A Study in Public Economy*, McGraw-Hill, New York; Oates, W.E. (1972), *Fiscal Federalism*, Harcourt Brace Jovanovich, New York; Oates, W.E. (1999), 'An Essay on Fiscal Federalism', *Journal of Economic Literature*, Vol. 37, No. 3, pp. 1120–1149, https://doi.org/10.1257/jel.37.3.1120.

Oberson, X. (2019), *Taxing Robots: Helping the Economy to Adapt to the Use of Artificial Intelligence*, Edward Elgar Publishing, Cheltenham; updated as Oberson, X. (2024), *Taxing Artificial Intelligence*, 2nd ed., Edward Elgar Publishing, Cheltenham, ISBN 9781035307548; and Oberson, X. (2025), 'Taxation of Artificial Intelligence', in Barfield, W. and Pagallo, U. (eds), *Research Handbook on the Law of Artificial Intelligence*, Edward Elgar Publishing, Cheltenham, pp. 636–651. The labour-market-impact residence criterion discussed in Section 5 was advanced by Oberson in response to Carvalho (see Carvalho, 2019).

OECD (2021), *Statement on a Two-Pillar Solution to Address the Tax Challenges Arising from the Digitalisation of the Economy*, OECD/G20 Inclusive Framework on BEPS, 8 October 2021, OECD, Paris (agreed by 136 member jurisdictions on the day; 139 as of 9 June 2023, per the OECD's own record), building on the Statement of 1 July 2021 (134 jurisdictions as of 31 August 2021).

OECD (2023a), *Multilateral Convention to Implement Amount A of Pillar One*, OECD/G20 Base Erosion and Profit Shifting Project, released 11 October 2023, OECD, Paris, https://www.oecd.org/tax/beps/multilateral-convention-to-implement-amount-a-of-pillar-one.pdf (in-scope thresholds: group revenue above €20 billion and profitability above 10%; reallocation of 25% of profit in excess of a 10% margin; the €20 billion threshold may fall to €10 billion seven years after entry into force, contingent on successful implementation).

OECD (2023b), *Outcome Statement on the Two-Pillar Solution to Address the Tax Challenges Arising from the Digitalisation of the Economy*, OECD/G20 Inclusive Framework on BEPS, 11 July 2023, approved by 138 members of the Inclusive Framework, OECD, Paris. Entry into force of the Multilateral Convention requires

ratification by 30 jurisdictions accounting for at least 60% of the ultimate parent entities of in-scope groups.

OECD (2025), *Tax Administration Digitalisation and Digital Transformation Initiatives*, op. cit., ch. 5 and Table 5.6.

OECD (2026a), *Consolidated Commentary to the Global Anti-Base Erosion Model Rules (2026)*, op. cit.; and OECD (2021), *Tax Challenges Arising from Digitalisation of the Economy – Global Anti-Base Erosion Model Rules (Pillar Two): Inclusive Framework on BEPS*, OECD Publishing, Paris, https://doi.org/10.1787/782bac33-en (Art. 5.3, substance-based income exclusion keyed to eligible payroll costs and eligible tangible assets, intangibles excluded from the asset base).

OECD (2026b), *Tax Challenges Arising from the Digitalisation of the Economy – Consolidated Commentary to the Global Anti-Base Erosion Model Rules (2026): Inclusive Framework on BEPS*, OECD/G20 Base Erosion and Profit Shifting Project, OECD Publishing, Paris, https://doi.org/10.1787/4377e89f-en, superseding OECD (2025), … *Consolidated Commentary … (2025)*, https://doi.org/10.1787/a551b351-en.

OECD et al. (2024), *Inventory of Tax Technology Initiatives*, OECD, Paris, https://www.oecd.org/en/data/datasets/inventory-of-tax-technology-initiatives.html accessed 7 February 2025), summarised in OECD (2025), *Tax Administration Digitalisation and Digital Transformation Initiatives*, OECD Publishing, Paris, https://doi.org/10.1787/c076d776-en, ch. 5 and Table 5.6 (72% of tax administrations use AI; the most common use case is the detection of tax evasion and fraud, undertaken by three quarters of them, the sub-percentages relating to the 72% that reported using AI).

Pasquale, F. (2015), *The Black Box Society: The Secret Algorithms That Control Money and Information*, Harvard University Press, Cambridge, MA, ISBN 9780674368279.

Ryngaert, C. (2015), *Jurisdiction in International Law*, 2nd ed., Oxford University Press, Oxford, https://doi.org/10.1093/law/9780199688517.001.0001 Ryngaert, C. (2020), *Selfless Intervention: The Exercise of Jurisdiction in the Common Interest*, Oxford University Press, Oxford.

Svantesson, D.J.B. (2017), *Solving the Internet Jurisdiction Puzzle*, Oxford University Press, Oxford, https://doi.org/10.1093/oso/9780198795674.001.0001.

Tomazela Santos, R. (2026), 'Pillar Two and the Tax Policy Challenges of the Substance-Based Income Exclusion', *Bulletin for International Taxation*, Vol. 80, No. 3, IBFD, Amsterdam, https://doi.org/10.59403/1pjarjk.

Turina, A. (2020), 'The progressive policy shift in the debate on the international tax challenges of the digital economy: a "pretext" for overhaul of the international tax regime?', *Computer Law & Security Review*, Vol. 36, art. 105382, https://doi.org/10.1016/j.clsr.2019.105382; see also Listokin, S. (2025), 'New rationales for taxing the digital economy: lessons from the OECD Pillar One consultations', *Frontiers in Political Science*, Vol. 7, art. 1561283, https://doi.org/10.3389/fpos.2025.1561283, reading that debate forward to artificial intelligence.



United Nations General Assembly (2023), *Promotion of Inclusive and Effective International Tax Cooperation at the United Nations*, Resolution A/RES/78/230, adopted 22 December 2023 at the 50th plenary meeting (111-46-10), establishing an ad hoc intergovernmental committee to draft terms of reference for a UN framework convention on international tax cooperation; and United Nations General Assembly (2024), Resolution A/RES/79/235, adopted 24 December 2024 (119-9-43), carrying the process forward.