

The Regulatory Stack as Trade Barrier: A GATS Analysis of the EU AI Act's Cumulative Impact with the GDPR on Cross-Border AI Services

Adam Dampe

Abstract

This paper argues that the EU Artificial Intelligence Act (Regulation (EU) 2024/1689), while largely origin-neutral on its face, creates — in combination with the GDPR — a systemic de facto barrier to cross-border AI services under the General Agreement on Trade in Services. Using the scenario of a US-based AI medical diagnostic service delivered from non-EU data centers to a public hospital consortium, the paper identifies two analytically distinct categories of AI Act-specific trade friction: cumulative regulatory burden (service-level regulation, deployer obligations, and conformity assessment) and a genuine regulatory stack interaction in which the AI Act's data governance requirements (Article 10) compound the GDPR's data transfer restrictions. Even under the EU-US Data Privacy Framework (Implementing Decision (EU) 2023/1795), the structural fragility of the adequacy mechanism leaves non-EU providers exposed to a compliance environment that their EU-based competitors do not face. While certain restrictions are justifiable under GATS Article XIV, the broader regulatory trajectory risks crossing from legitimate regulatory autonomy into unjustifiable trade restriction — a tension that the ongoing crisis of the WTO Appellate Body leaves without a clear path to resolution.

The Regulatory Stack as Trade Barrier: A GATS Analysis of the EU AI Act's Cumulative Impact with the GDPR on Cross-Border AI Services

Adam Dampe University of Heidelberg April 2026

Abstract

Background: The EU Artificial Intelligence Act (Regulation (EU) 2024/1689) is the world's first comprehensive binding regulation of AI systems. While a growing body of scholarship examines its compatibility with WTO law, existing analyses focus primarily on the AI Act in isolation and on its prohibited practices.

The cumulative trade impact of the AI Act layered on top of the GDPR's data transfer restrictions remains underexplored.

Method: This paper conducts a GATS-focused analysis of the AI Act's regulatory architecture, applying the non-discrimination tests (MFN, National Treatment) and the general exceptions framework (Article XIV) to the combined regulatory stack of the AI Act and the GDPR. It uses the scenario of a US-based AI medical diagnostic service supplied cross-border from non-EU data centers to a public hospital consortium within the EU.

Findings: The AI Act's formal provisions are largely origin-neutral and survive direct non-discrimination scrutiny. However, the cumulative regulatory architecture creates two analytically distinct categories of de facto trade friction for cross-border AI services. First, the AI Act imposes cumulative regulatory burden through service-level requirements, deployer obligations, and conformity assessment procedures that go beyond the GDPR's data-level restrictions. Second — and more significantly — the AI Act's data governance requirements (Article 10) interact with the GDPR's data transfer restrictions to produce a genuine regulatory stack effect: a compounding dynamic in which compliance with one regulation's requirements is made structurally harder by the other's constraints. Even under the EU-US Data Privacy Framework, the structural fragility of adequacy mechanisms means this interaction disproportionately affects non-EU providers. While prohibitions such as social scoring are justifiable under Article XIV(a) (public order), the broader regulatory trajectory — particularly enforcement discretion and standardization processes — could cross from legitimate regulatory autonomy into unjustifiable trade restriction.

Implications: With the WTO Appellate Body non-functional since 2019, the tension between the EU's regulatory ambition and its GATS obligations lacks a clear adjudicative resolution path. Administrative practice under the AI Act should be monitored for de facto discriminatory effects, and the cumulative burden of the EU's regulatory stack deserves sustained attention from both trade and technology law scholars.

Keywords: EU AI Act, GATS, trade barriers, GDPR, cross-border AI services, de facto discrimination, regulatory stack, WTO dispute settlement, Article XIV GATS general exceptions, AI regulation and international trade, EU-US Data Privacy Framework

1. Introduction

The Artificial Intelligence Act (Regulation (EU) 2024/1689) entered into force on 1 August 2024, establishing the world's first comprehensive binding regulatory framework for AI systems. Its risk-based classification scheme, conformity assessment requirements, and extensive deployer obligations represent the most ambitious attempt by any jurisdiction to subject AI systems to systematic legal

governance. The regulation is explicitly extraterritorial in scope: it applies to providers and deployers regardless of whether they are established within the EU, as long as their AI systems are placed on the EU market or their outputs affect persons located in the EU. The AI Act is, in this sense, a regulatory manifestation of what Bradford (2020) has termed the “Brussels Effect” — the EU’s capacity to shape global regulatory standards through the sheer gravitational pull of its internal market. This paper does not examine the political economy of that dynamic; it examines its trade law consequences.

A growing body of scholarship has begun to examine the AI Act’s compatibility with the multilateral trade rules of the World Trade Organization (WTO). Most notably, Soprana (2024) analyzed the compatibility of the AI Act’s prohibited AI systems with both the Agreement on Technical Barriers to Trade (TBT) and the General Agreement on Trade in Services (GATS), concluding that there is potential for conflict between the EU regulation and both multilateral agreements. Other contributions have examined AI regulation and international trade at a higher level of generality, addressing the classification of AI under the goods-services dichotomy and the challenges that AI poses for WTO governance structures (Liu and Lin 2020; WTO 2024).

This paper departs from the existing literature in two respects. First, it focuses exclusively on GATS rather than splitting the analysis between TBT and GATS, allowing for a more granular examination of the services trade implications. Second — and more significantly — it does not analyze the AI Act in isolation. Instead, it examines the cumulative trade impact of the AI Act layered on top of the General Data Protection Regulation (GDPR, Regulation (EU) 2016/679). The central argument is that neither regulation alone triggers a clear GATS violation, but their interaction can create a regulatory stack that produces de facto barriers to cross-border AI services. The concept of cumulative regulatory burden is not new — scholarship on digital trade fragmentation (Burri 2021) and data protection as trade barrier (Aaronson 2018; Selby 2017) has examined similar dynamics. What distinguishes the regulatory stack argument developed here is its focus on architectural interdependence: how one regulation’s requirements amplify the restrictions imposed by another, producing non-linear compliance effects that neither creates independently.

The analysis is grounded in a concrete scenario: a US-based provider of an AI-powered medical diagnostic service, delivered cross-border (GATS Mode 1) from data centers located outside the EU to a public hospital consortium within the EU. This is deliberately a “hard case” — it maximizes the trade law tension by combining a high-risk AI classification, non-EU data infrastructure, and the full weight of both GDPR and AI Act requirements. The argument weakens significantly for low-risk AI services (which face minimal conformity assessment and no deployer FRIA), for providers from jurisdictions with GDPR adequacy decisions (where data transfer restrictions are attenuated), and for providers with EU commercial presence (Mode 3, which simplifies regulatory navigation). Acknowledging these boundary conditions strengthens rather than undermines

the analysis: it shows precisely where the regulatory stack becomes problematic and where it does not.

The paper proceeds as follows. Section 2 maps the AI Act’s regulatory architecture through a trade lens. Section 3 establishes the GATS framework for AI services. Section 4 applies the non-discrimination tests — MFN (Article II) and National Treatment (Article XVII). Section 5 examines the cumulative regulatory burden that the AI Act adds beyond the GDPR. Section 6 develops the paper’s core original contribution: the regulatory stack interaction between Article 10 AI Act data governance requirements and GDPR data transfer restrictions. Section 7 analyzes justification under GATS Article XIV general exceptions. Section 8 concludes with an assessment of the regulatory trajectory in light of the WTO Appellate Body crisis.

2. The AI Act’s Regulatory Architecture Through a Trade Lens

The AI Act establishes a risk-based classification scheme under which AI systems are categorized as posing unacceptable, high, limited, or minimal risk, with regulatory obligations scaled accordingly. For the purposes of trade law analysis, four elements of the regulatory architecture are particularly significant.

2.1 Prohibited Practices and High-Risk Classification

Article 5 of the AI Act prohibits certain AI practices outright, including social scoring systems and real-time remote biometric identification in public spaces (with limited exceptions for law enforcement). These prohibitions apply regardless of the origin of the AI system or its provider. For a non-EU provider, a system classified as prohibited cannot be placed on the EU market at all — this amounts to a complete market access barrier for specific types of AI services.

The high-risk classification (Article 6, Annex III) encompasses AI systems deployed in eight areas including biometrics, critical infrastructure, education, employment, access to essential services, law enforcement, migration, and administration of justice. A US-based AI medical diagnostic tool falls within this category as a system intended for use as a medical device. High-risk classification triggers the full suite of compliance obligations discussed below.

2.2 Provider Obligations and the Value Chain

Articles 9–15 impose requirements on providers of high-risk AI systems covering risk management (Article 9), data governance (Article 10), technical documentation (Article 11), record-keeping (Article 12), transparency to deployers (Article 13), human oversight design (Article 14), and accuracy, robustness, and cybersecurity (Article 15). These are requirements on the system as a product, not merely on its data processing activities — a distinction that becomes important in Section 5.

Article 25 distributes responsibilities along the AI value chain between providers, deployers, distributors, importers, and authorized representatives. This distribution partially mitigates the compliance burden on any single actor, but it also means that cross-border supply chains face coordination costs that purely domestic arrangements avoid.

2.3 Deployer Obligations

Article 26 imposes obligations on deployers of high-risk AI systems, including ensuring human oversight, monitoring for risks, maintaining logs, and informing affected persons. Article 27 requires a Fundamental Rights Impact Assessment (FRIA) before deploying certain high-risk AI systems — but this obligation is narrower than sometimes assumed: it applies to “deployers that are bodies governed by public law, or . . . private entities providing public services” and to specific high-risk categories involving biometrics (Annex III, point 1) and access to essential services (Annex III, point 6). Private entities not providing public services are not automatically subject to the FRIA requirement.

For the MedAI scenario, a public hospital consortium deploying an AI diagnostic tool would fall within Article 27’s scope. The FRIA must be conducted with reference to fundamental rights as understood under the EU Charter of Fundamental Rights — a framework reflecting distinctly European normative commitments regarding dignity, non-discrimination, and data protection that may differ from the constitutional frameworks of other major AI-exporting jurisdictions.

2.4 Conformity Assessment and Standardization

Articles 40–43 establish conformity assessment procedures for high-risk AI systems. For medical devices incorporating AI, this generally involves third-party assessment by notified bodies designated by EU Member States (Article 43, read with Annex VII). Medical AI systems are already subject to conformity assessment under the Medical Devices Regulation (MDR, Regulation (EU) 2017/745); the AI Act adds a layer to this existing framework rather than creating a wholly new mechanism. Nevertheless, the additional AI-specific requirements — particularly the data governance and transparency obligations — represent incremental compliance cost.

The standardization process raises its own trade-relevant concerns. Harmonized standards under the AI Act are developed by European standardization organizations (CEN, CENELEC) with predominantly European industry participation. While non-EU stakeholders can participate, practical barriers — time zones, language, institutional familiarity, and the cost of sustained engagement — may mean that the standard-setting process disproportionately reflects European industry practices. Article 57 further provides for AI regulatory sandboxes that allow supervised market entry; while formally open to all eligible applicants regardless of origin (Article 57(1), Recital 138), the operational reality of sustained engagement with national authorities may favour proximity.

3. GATS Framework for AI Services

3.1 Classification of AI Under GATS

For cloud-based AI services delivered cross-border — the dominant commercial model for AI-as-a-service offerings including medical diagnostics — classification as a service under GATS is straightforward. The AI system is not a tangible product shipped across borders; it is a computational capability accessed remotely. This falls within GATS Mode 1 (cross-border supply), where the service crosses the border while neither provider nor consumer moves. The Mode 1 classification is appropriate here because MedAI operates its infrastructure outside the EU and supplies the diagnostic capability remotely; there is no commercial presence (Mode 3) and no movement of natural persons (Mode 4). This classification matters because Mode 1 is generally the most liberalized mode in the EU’s schedule, but it is also the mode in which “commercial presence” advantages become most relevant to the de facto discrimination analysis.

The specific GATS subsector classification is less clear. An AI medical diagnostic service could be classified under computer and related services (CPC 84), health-related services (CPC 93), or professional services. For computer and related services, the EU has generally undertaken relatively liberal commitments; for health services, commitments are more limited. The classification question itself introduces legal uncertainty that could affect the trade law analysis.

3.2 The Running Scenario

Consider MedAI Inc., a US-based company that has developed an AI-powered radiological diagnostic tool. The system analyzes medical imaging data (X-rays, CT scans, MRI) and provides diagnostic suggestions to clinicians. MedAI operates its computing infrastructure from data centers in Virginia and Oregon. A consortium of public hospitals in Germany wishes to procure MedAI’s service for clinical decision support.

Under the GDPR alone, MedAI faces significant but manageable hurdles: it must establish a lawful basis for processing EU patients’ medical imaging data (health data under Article 9 GDPR), implement appropriate safeguards for international data transfers (Chapter V GDPR), and comply with data subject rights. Since the EU-US Data Privacy Framework (DPF) adequacy decision of July 2023 (Implementing Decision (EU) 2023/1795), MedAI can — if certified under the DPF — lawfully transfer EU health data to its US data centers without Standard Contractual Clauses or Binding Corporate Rules, subject to the DPF’s enhanced safeguards. The significance and fragility of this mechanism is addressed in Section 6.

Under the AI Act layered on top of the GDPR, MedAI faces an additional set of obligations: its system must be classified and registered as high-risk (Article 6, Annex III), subjected to conformity assessment potentially involving a European notified body (Article 43), equipped with comprehensive technical documentation

(Article 11), designed for transparency to deployers (Article 13), built with human oversight capabilities (Article 14), and subject to ongoing post-market monitoring (Article 72). Meanwhile, the deploying hospital consortium — as bodies governed by public law — must conduct a FRIA (Article 27), implement human oversight (Article 26), maintain logs, and inform patients that AI is being used in their diagnostic process.

This is deliberately a hard case. As noted in Section 1, the argument weakens considerably for low-risk AI services, for providers from GDPR-adequate jurisdictions (such as the UK, Japan, or South Korea, where the data transfer dimension is attenuated), and for providers that establish EU commercial presence. The hard case is analytically useful precisely because it reveals the boundary conditions of GATS compliance — the point at which the regulatory stack becomes most problematic.

4. The Non-Discrimination Analysis

4.1 Most-Favoured-Nation Treatment (Article II GATS)

Article II:1 GATS requires that each WTO Member accord to services and service suppliers of any other Member treatment no less favourable than that it accords to like services and service suppliers of any other country. The AI Act does not formally differentiate between AI services or providers from different WTO Members. A Chinese AI provider and a US AI provider face identical regulatory requirements. On its face, the AI Act does not violate MFN.

However, the Appellate Body in EC — Bananas III (WT/DS27/AB/R) established that the MFN obligation under GATS applies to both de jure and de facto discrimination. It should be noted that the Appellate Body's finding in that case rested primarily on explicit origin-based licensing distinctions (the “operator categories”) rather than purely de facto effects, and the de facto dimension was not the primary ratio decidendi. Nevertheless, the principle that Article II covers de facto discrimination was affirmed. The Appellate Body in Argentina — Financial Services (WT/DS453/AB/R) further clarified that formally origin-neutral measures can violate GATS non-discrimination obligations if they modify the conditions of competition to the detriment of like services or service suppliers of another Member. Critically, the Appellate Body rejected the argument that “regulatory aspects” could justify such modifications within the non-discrimination analysis itself — regulatory justifications are appropriately addressed under Article XIV, not as a defense within the non-discrimination test.

Applied to the AI Act: if administrative practices under the regulation — enforcement discretion by national market surveillance authorities, the speed of conformity assessment procedures, or the practical accessibility of regulatory sandboxes — were to systematically disadvantage providers from specific third countries, this could constitute de facto MFN discrimination. The formal neutrality of the regulation does not insulate actual enforcement practice from scrutiny.

However, this remains a forward-looking risk rather than a demonstrated violation — the AI Act’s high-risk provisions only become fully applicable in August 2026, and enforcement practice has yet to develop a track record.

4.2 National Treatment (Article XVII GATS)

Article XVII requires that, in sectors where specific commitments are inscribed, each Member accord to services and service suppliers of other Members treatment no less favourable than its own like services and service suppliers. The AI Act appears formally neutral, but the national treatment analysis must consider whether formally identical treatment produces *de facto* less favourable conditions for foreign service suppliers.

Two structural features of the AI Act could produce such effects. First, conformity assessment involving EU-based notified bodies is procedurally easier for providers with established EU presence — they are familiar with European standardization processes and can engage in the regulatory culture of the assessment process. However, this advantage stems in part from incumbency rather than from the AI Act itself; a new EU-based provider entering the market for the first time would face similar learning costs, though without the additional complexity of cross-border coordination. Second, the cumulative weight of AI Act requirements layered on top of GDPR compliance (developed in Sections 5 and 6) may produce conditions of competition that structurally favour EU-based providers, particularly those operating EU data infrastructure.

The *de facto* national treatment analysis under Argentina — Financial Services requires demonstrating that the measures modify the conditions of competition to the detriment of foreign services or service suppliers. This threshold is significant, and the paper acknowledges that the evidence base for competitive effects under the AI Act remains limited — the regulation’s high-risk provisions are not yet fully in force, and empirical data on compliance costs for EU versus non-EU providers does not yet exist. The argument developed here identifies mechanisms through which *de facto* discrimination could arise, not established facts of discrimination. As Krajewski (2003) has shown, the scope of national treatment in services trade extends well beyond explicitly discriminatory measures, encompassing regulatory frameworks that structurally disadvantage foreign service suppliers — a framing that informs the analysis developed in Sections 5 and 6.

5. Cumulative Regulatory Burden Beyond the GDPR

Before turning to the paper’s core contribution — the regulatory stack interaction — it is necessary to establish what the AI Act adds to the GDPR in terms of straightforward cumulative burden. Three layers of additional regulatory friction apply to cross-border AI services regardless of their interaction with GDPR data transfer restrictions.

5.1 Service-Level Regulation Beyond Data-Level Restriction

The GDPR governs how personal data is processed. It does not regulate the service itself — it does not prescribe how an AI diagnostic tool must be designed, what level of accuracy it must achieve, how transparent it must be, or what risk management processes its provider must maintain. The AI Act does all of these things. Articles 9–15 impose requirements on the system as a product: risk management, data governance, technical documentation, record-keeping, transparency, human oversight design, and accuracy. For a cross-border AI service provider, GDPR compliance is necessary but not sufficient. Full compliance with every GDPR requirement still leaves the AI Act’s service-level requirements entirely unaddressed. Under GATS, this represents an additional layer of conditions on the supply of services that did not exist before August 2024.

5.2 Demand-Side Deployer Burden

Under the GDPR, the primary compliance burden falls on the data controller or processor — typically the service provider or the entity determining the purposes of processing. Under the AI Act, substantial obligations shift to the deployer. As described in Section 2.3, deployer obligations include human oversight, monitoring, logging, transparency to affected persons, and — for public bodies and public service providers — the FRIA. Article 25 distributes responsibilities across the value chain, partially mitigating the burden on deployers, but the net effect remains that the EU entity procuring a foreign AI service bears regulatory costs that do not arise for non-AI alternatives. A public hospital that employs radiologists faces no AI Act obligations; the same hospital deploying MedAI’s diagnostic tool faces the full suite of deployer requirements. This demand-side friction dampens procurement of AI services generally and, in practice, may disproportionately affect foreign providers who cannot offer the same level of hands-on compliance support as locally established firms.

5.3 Conformity Assessment as Market Access Mechanism

The AI Act’s conformity assessment regime introduces a pre-market approval mechanism for high-risk AI systems that the GDPR does not require. The GDPR relies on ex post enforcement through supervisory authorities; the AI Act imposes ex ante certification. A cross-border AI service cannot be lawfully supplied to EU deployers until it has been certified as compliant — a market access condition under GATS Article XVI. For MedAI, this means undergoing a separate conformity assessment process (in addition to any MDR assessment for medical devices) involving European notified bodies, European harmonized standards, and European quality management system oversight.

These three layers — service-level regulation, deployer burden, and conformity assessment — are additive. They would apply even if the GDPR did not exist, and they would burden a purely domestic EU provider entering the market for the first

time as well (though with lower transaction costs due to institutional proximity). They represent genuine regulatory burden, but they do not, standing alone, make the regulatory stack argument. That argument requires demonstrating that the AI Act and GDPR interact to produce effects that neither creates independently. This is the task of Section 6.

6. The Regulatory Stack Interaction: Article 10 Data Governance Meets GDPR Data Transfers

This section develops the paper’s core original contribution.

6.1 The Architecture of the Interaction

Article 10 of the AI Act requires that training, validation, and testing datasets for high-risk AI systems be “relevant, sufficiently representative, and to the best extent possible, free of errors and complete in view of the intended purpose” (Article 10(3)). Datasets must “take into account, to the extent required by the intended purpose, the characteristics or elements that are particular to the specific geographical, contextual, behavioural or functional setting within which the high-risk AI system is intended to be used” (Article 10(4)). For bias detection and correction, providers may “exceptionally” process special categories of personal data under strict conditions (Article 10(5)).

Read in isolation, these are reasonable data quality requirements. Read together with the GDPR’s restrictions on international data transfers, they create a compounding dynamic. To develop an AI diagnostic tool that is “sufficiently representative” of the EU population and accounts for the “geographical... setting” of EU deployment, MedAI needs access to datasets that reflect the demographic, clinical, and epidemiological characteristics of EU patient populations. The GDPR restricts precisely this: transferring EU patients’ health data (a special category under Article 9 GDPR) to data centers in Virginia requires robust legal mechanisms under Chapter V GDPR.

It should be acknowledged that Article 10(4)’s representativeness requirement can be partially satisfied through validation and testing datasets constructed within the EU, without requiring retraining of the underlying model on EU data. Not all providers retrain their models at deployment. However, for AI systems intended for high-risk medical use, ongoing representativeness — particularly across diverse EU populations — typically requires iterative refinement with local data. A provider that cannot access EU patient data for any model improvement operates at a structural disadvantage compared to a provider with EU data center infrastructure.

6.2 The EU-US Data Privacy Framework and Its Fragility

The analysis would be incomplete without addressing the EU-US Data Privacy Framework (DPF), established by Commission Implementing Decision (EU)

2023/1795 of 10 July 2023. The DPF provides an adequacy mechanism that allows DPF-certified US entities to receive EU personal data — including health data — without relying on Standard Contractual Clauses or Binding Corporate Rules.

The DPF significantly attenuates the compounding trap described above. A DPF-certified US AI provider can lawfully receive EU health data for model development and validation, provided it complies with the DPF principles (including purpose limitation, data minimization, and the redress mechanism established under US Executive Order 14086). To this extent, the regulatory stack interaction is currently less acute than it would be in the absence of an adequacy framework.

However, the DPF does not eliminate the regulatory stack interaction for three reasons.

First, the DPF's legal foundation is structurally fragile. It rests on US Executive Order 14086, which can be modified or revoked by a future US administration without congressional approval. The Court of Justice of the European Union has already invalidated two predecessor frameworks — Safe Harbor in *Schrems I* (Case C-362/14) and the Privacy Shield in *Schrems II* (Case C-311/18). An annulment action brought by French Member of Parliament Philippe Latombe before the General Court was dismissed on 3 September 2025 (Case T-553/23), but this dismissal does not resolve the substantive question of the DPF's durability — privacy advocacy organization NOYB has indicated it is examining the Latombe judgment and the possibility of further challenges remains open. A non-EU AI provider that builds its business model on the DPF faces an adequacy mechanism whose long-term stability cannot be assured.

Second, DPF certification is voluntary and far from universal. Not all US AI companies — particularly smaller or specialized providers — have undergone DPF certification. For non-certified providers, the pre-DPF landscape of Standard Contractual Clauses with Transfer Impact Assessments continues to apply, preserving much of the compounding effect.

Third, even under the DPF, the asymmetry between EU-based and US-based providers persists. An EU provider operating EU data centers faces no data transfer question at all — it processes EU health data domestically. A US provider, even with DPF certification, must navigate a transfer mechanism with ongoing compliance obligations, potential future invalidation, and the reputational risk associated with cross-border health data flows. The AI Act's Article 10 requirements — which demand that data reflect the specific EU deployment context — amplify this asymmetry by raising the stakes of data access. The AI Act does not merely require data processing (which the GDPR governs); it requires representative data processing (which depends on accessing data that reflects local conditions). This interaction effect is the crux of the regulatory stack argument.

6.3 Structural Advantage, Not Formal Discrimination

The compounding dynamic does not formally discriminate by origin. An EU provider that chose to operate from US data centers would face the same regulatory stack. But the structural reality is that EU-based AI providers operating EU data infrastructure face no transfer mechanism requirement for accessing EU training and validation data, while non-EU providers — even DPF-certified ones — do. Under the de facto discrimination analysis developed in Argentina — Financial Services, such structural effects on the conditions of competition are relevant to the national treatment assessment. Whether they rise to the level of “modifying the conditions of competition to the detriment of like services or service suppliers of another Member” is ultimately a question that would require adjudication — adjudication that, as discussed in Section 8, is currently unavailable.

7. Justification Under GATS Article XIV

Assuming that the regulatory stack could create de facto barriers to cross-border AI services, the question becomes whether these barriers can be justified under the GATS general exceptions.

7.1 Public Morals and Public Order (Article XIV(a))

Article XIV(a) allows Members to adopt measures “necessary to protect public morals or to maintain public order.” The public order exception carries a heightened threshold: the measure may be invoked “only where a genuine and sufficiently serious threat is posed to one of the fundamental interests of society” (Footnote 5 to Article XIV).

For the AI Act’s prohibited practices — notably social scoring and certain biometric surveillance applications — the Article XIV(a) justification is strong. These prohibitions reflect the EU’s fundamental rights commitments and address what the EU considers genuine threats to democratic society. The FRIA requirement, rooted in the EU Charter of Fundamental Rights, similarly reflects core public order concerns. The EU’s position here is not merely a regulatory preference but a constitutional constraint: the EU Charter is primary law, and GATS does not require Members to abandon constitutional protections for commercial convenience.

For the broader high-risk regime, the public order justification is less straightforward. The protection of health and safety through medical device regulation is unquestionably a legitimate objective, but the necessity test requires demonstrating that the specific measures are not more trade-restrictive than necessary to achieve the regulatory purpose.

7.2 The Necessity Test and the US — Gambling Precedent

The necessity test under Article XIV, developed through GATT Article XX jurisprudence and imported into the GATS context by the Appellate Body in US — Gambling (WT/DS285/AB/R), requires a weighing and balancing of the importance of the regulatory objective, the contribution of the measure, and its trade-restrictiveness. Critically, it also requires consideration of reasonably available less trade-restrictive alternatives.

The US — Gambling precedent creates a tension that this paper must acknowledge. In that case, the United States lost its Article XIV defense in part because the Appellate Body found that less trade-restrictive alternatives — including Antigua’s offer of regulatory cooperation — were reasonably available. This precedent cuts against an easy Article XIV justification for the AI Act’s broader requirements. If a WTO Member could demonstrate that mutual recognition of conformity assessments from equivalent third-country regulatory systems, or regulatory cooperation mechanisms allowing supervised market entry without full EU conformity assessment, were reasonably available alternatives, the EU’s necessity defense would face significant difficulty.

The record of Article XIV defenses more broadly is mixed. Only a small number of attempts to invoke GATT Article XX or GATS Article XIV general exceptions have succeeded, and many have failed at the chapeau stage — precisely the stage where discriminatory application (as opposed to discriminatory design) is assessed. The chapeau requires that measures not be applied in a manner constituting “arbitrary or unjustifiable discrimination between countries where like conditions prevail” — as demonstrated in US — Shrimp and Brazil — Retreaded Tyres, even facially neutral measures frequently fail this test due to discriminatory implementation.

7.3 The Chapeau and Administrative Practice

This is where the paper’s warning becomes most concrete. The AI Act as written may well survive the Article XIV analysis: its objectives (protection of health, safety, and fundamental rights) are among the most clearly recognized under GATS, and its provisions are formally origin-neutral. But the regulatory trajectory — the development of enforcement practice, the elaboration of harmonized standards, the exercise of discretion by the AI Office and national market surveillance authorities — creates ongoing risk.

If enforcement authorities apply conformity assessment procedures more slowly or stringently to non-EU providers; if harmonized standards embed assumptions reflecting European industry practices without adequate international input; if the cumulative interaction between AI Act and GDPR requirements creates structural advantages for providers with EU data center presence — then the formally neutral regulatory framework could produce effects that would not survive the chapeau analysis. The key constraint is this: GATS non-discrimination obligations mean that US, Chinese, or other third-country AI providers cannot be

treated differently on a country-by-country basis through enforcement discretion, unless such differentiation is independently justified under the general exceptions. Administrative practice must respect this boundary.

8. Conclusion

The EU AI Act is not, as of this writing, a demonstrated trade barrier. Its formal provisions are largely origin-neutral, and its regulatory objectives — protection of health, safety, and fundamental rights — are among the most clearly legitimate under GATS general exceptions. The classic non-discrimination tests (MFN under Article II, National Treatment under Article XVII) do not bite directly on the AI Act’s text.

However, the AI Act does not operate in a vacuum. When layered on top of the GDPR, it creates two analytically distinct categories of trade friction for cross-border AI services. The first is cumulative regulatory burden: service-level requirements beyond data-level restrictions, deployer obligations that dampen procurement, and conformity assessment as a pre-market access condition. The second — and more significant — is the regulatory stack interaction between Article 10 data governance requirements and GDPR data transfer restrictions. This interaction creates a structural asymmetry between EU-based and non-EU providers that persists even under the Data Privacy Framework, because the DPF attenuates but does not eliminate the data access differential, and because its own legal foundation remains structurally fragile.

The non-discrimination analysis, informed by the *de facto* discrimination doctrine of EC — *Bananas III* and *Argentina — Financial Services*, suggests that these cumulative effects could be cognizable under GATS — though whether they would ultimately be found to “modify the conditions of competition” to the required degree is a question that would require adjudication on the basis of empirical evidence that does not yet exist. Whether potential justifications under Article XIV would succeed depends critically on implementation choices: the US — *Gambling* precedent demonstrates that even legitimate regulatory objectives can fail the necessity test when less trade-restrictive alternatives — such as mutual recognition or regulatory cooperation — are reasonably available.

The situation is further complicated by the ongoing crisis of the WTO Appellate Body. Since the term of its last sitting member expired in November 2020, the Appellate Body has been unable to hear any appeals. Reform efforts have stalled — despite a commitment at the 2022 Ministerial Conference to restore a fully functional dispute settlement system by 2024, that deadline was missed, and most panel reports are now appealed “into the void.” The Multi-Party Interim Appeal Arbitration Arrangement provides a partial substitute, but it lacks the participation of major trading nations including the United States. With no clear resolution in sight, the AI Act’s GATS compatibility cannot be authoritatively tested through the WTO’s own dispute settlement machinery.

This paper therefore serves both as a doctrinal contribution and a forward-

looking analysis. The doctrinal contribution is the regulatory stack argument: the demonstration that layered regulations can interact to produce trade barriers that neither creates independently. The forward-looking analysis is that the EU's implementation of the AI Act — the choices made by the AI Office, by national authorities, by standardization bodies — will determine whether the regulatory framework remains within the bounds of the EU's international trade commitments. Whether the AI Act becomes a model of legitimate regulatory autonomy or a case study in how regulatory accumulation can produce unjustifiable trade effects depends on implementation choices not yet made. The burden of that question will be particularly acute for AI providers from developing countries that lack the regulatory infrastructure, institutional relationships, and financial resources to navigate the EU's layered compliance landscape — a dimension that warrants sustained attention in future research.

Bibliography

- Aaronson, Susan Ariel (2018). “Data Is Different: Why the World Can't Agree on How to Govern Data and Cross-Border Data Flows.” CIGI Paper No. 197. Waterloo: Centre for International Governance Innovation. URL: https://www.cigionline.org/static/documents/documents/paper%20no.197_0.pdf
- Bradford, Anu (2020). *The Brussels Effect: How the European Union Rules the World*. Oxford: Oxford University Press. DOI: <https://doi.org/10.1093/oso/9780190088583.001.0001>
- Burri, Mira (2021). “Data Flows and Global Trade Law.” In: Burri, Mira (ed.), *Big Data and Global Trade Law*. Cambridge: Cambridge University Press, pp. 11–41. DOI: <https://doi.org/10.1017/9781108919234.003>
- Cottier, Thomas, Delimatsis, Panagiotis, and Diebold, Nicolas F. (2008). “Article XIV GATS: General Exceptions.” In: Wolfrum, Rüdiger, Stoll, Peter-Tobias, and Feinäugle, Clemens (eds.), *Max Planck Commentaries on World Trade Law, WTO — Trade in Services, Vol. 6*, pp. 287–328. Leiden/Boston: Martinus Nijhoff Publishers. URL: <https://ssrn.com/abstract=1280215>
- Delimatsis, Panagiotis, and Gargne, Léo (2020). “General Exceptions under the GATS — A Legal Commentary on Article XIV GATS.” TILEC Discussion Paper Series No. 2020-027, pp. 1–39. URL: <https://ssrn.com/abstract=3757464>
- Muller, Gilles (2017). “De facto Discrimination Under GATS National Treatment: Has the Genie of Trade Liberalization Been Let Out of the Bottle?” *Legal Issues of Economic Integration* 44(2): 151–172. DOI: <https://doi.org/10.54648/LEIE2017009>
- Liu, Han-Wei, and Lin, Ching-Fu (2020). “Artificial Intelligence and Global Trade Governance: A Pluralist Agenda.” *Harvard International Law Journal* 61(2). URL: <https://journals.law.harvard.edu/ilj/wp-content/uploads/sites/84/61.2-Liu.pdf>

Krajewski, Markus (2003). National Regulation and Trade Liberalization in Services: The Legal Impact of the General Agreement on Trade in Services (GATS) on National Regulatory Autonomy. The Hague: Kluwer Law International. [no DOI/URL available]

Selby, John (2017). “Data Localization Laws: Trade Barriers or Legitimate Responses to Cybersecurity Risks, or Both?” *International Journal of Law and Information Technology* 25(3): 213–232. DOI: <https://doi.org/10.1093/ijlit/eax010>

Soprana, Marta (2024). “Compatibility of Emerging AI Regulation with GATS and TBT: The EU Artificial Intelligence Act.” *Journal of International Economic Law* 27(4): 706–722. DOI: <https://doi.org/10.1093/jiel/jgae040>

WTO Appellate Body (1997). European Communities — Regime for the Importation, Sale and Distribution of Bananas (EC — Bananas III). WT/DS27/AB/R. URL: https://www.wto.org/english/tratop_e/dispu_e/cases_e/ds27_e.htm

WTO Panel and Appellate Body (2005). United States — Measures Affecting the Cross-Border Supply of Gambling and Betting Services (US — Gambling). WT/DS285/AB/R. URL: https://www.wto.org/english/tratop_e/dispu_e/cases_e/ds285_e.htm

WTO Panel and Appellate Body (2016). Argentina — Measures Relating to Trade in Goods and Services (Argentina — Financial Services). WT/DS453/AB/R. URL: https://www.wto.org/english/tratop_e/dispu_e/cases_e/ds453_e.htm

World Trade Organization (2024). *Trading with Intelligence: How AI Shapes and Is Shaped by International Trade*. Geneva: WTO Publications. URL: https://www.wto.org/english/res_e/booksp_e/trading_with_intelligence_e.pdf

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). OJ L 2024/1689. URL: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). OJ L 119/1. URL: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>

Commission Implementing Decision (EU) 2023/1795 of 10 July 2023 pursuant to Regulation (EU) 2016/679 of the European Parliament and of the Council on the adequate level of protection of personal data under the EU-US Data Privacy Framework. OJ L 231/118. URL: https://eur-lex.europa.eu/eli/dec_impl/2023/1795/oj

Court of Justice of the European Union (2020). *Data Protection Commissioner v. Facebook Ireland Limited and Maximilian Schrems (Schrems II)*. Case C-311/18, ECLI:EU:C:2020:559. URL: <https://curia.europa.eu/juris/liste.jsf?num=C-311/18>

Court of Justice of the European Union (2015). *Maximilian Schrems v. Data*

Protection Commissioner (Schrems I). Case C-362/14, ECLI:EU:C:2015:650.
URL: <https://curia.europa.eu/juris/liste.jsf?num=C-362/14>
