

MASTER THESIS | MASTER'S THESIS

Titel | Title

“The Brussels Effect and the Limits of Artificial Intelligence Governance: The AI Act in the U.S. constitutional order”

verfasst von | submitted by

Katharina Gmeiner, licence

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of
Master of Advanced International Studies (M.A.I.S.)

Wien | Vienna, 2026

Studienkennzahl lt. Studienblatt | UA 992 940
Degree programme code as it appears
on the student record sheet:

Universitätslehrgang lt. Studienblatt | Internationale Studien | International Studies
Degree programme as it appears on
the student record sheet:

Betreut von | Supervisor: Professor Dr. Holger Hestermeyer



**diplomatische
akademie wien**

Vienna School of International Studies
École des Hautes Études Internationales de Vienne

Pledge of Honesty and Declaration of AI Assistance

Katharina Gmeiner

On my honour as a student of the Diplomatische Akademie Wien, I submit this work in good faith and pledge that I have neither given nor received unauthorized assistance on it.

Sources for literal, paraphrased and cited quotes have been accurately credited. I hereby declare that this Master's thesis was independently written and researched by me. All contents, arguments, analysis and conclusions are my own. While completing this thesis, I made use of AI-assisted tools for the following purposes: grammar, assistance with chapter structuring and outline development. In all cases, the final decisions regarding content, structure, and academic argumentation were made solely by me. AI was used strictly as a supportive tool and did not generate original research content or conclusions. AI tools used for this study are listed below.

This declaration is made in the interest of transparency and academic integrity. The submitted document here presented is identical to the electronically submitted text document.

AI tools that I have used:

- *Open AI – ChatGPT*
- *Anthropic – Claude AI*

Abstract:

AI systems are deployed globally through centralized architectures that serve users across jurisdictions simultaneously. The European Union's (EU) Artificial Intelligence Act (AI Act), the first comprehensive binding regulatory framework for AI, responds to this by asserting jurisdiction over any provider whose systems affect persons within the Union, regardless of the establishment location. The AI Act operationalizes this through a market access logic: any provider wishing to serve users within the Union must comply, regardless of where that provider is established. The theoretical consequences of the Brussels Effect are clear: firms, including firms from the United States of America (U.S.), choose to extend EU compliance globally rather than fragment their operations, allowing the EU to shape AI governance beyond its borders through market pressure alone.

When U.S. firms apply EU AI Act obligations globally and therefore implement EU compliance standards across their entire user base, those decisions produce effects on First Amendment-protected speech that U.S. constitutional doctrine is structurally positioned to evaluate but only where a cognizable legal claim can be brought. Courts do not apply EU law. What they review is domestic conduct, which was shaped by foreign compliance incentives, measured against settled domestic constitutional standards. Were Congress to enact equivalent obligations, they would raise serious First Amendment problems. U.S. and EU free speech frameworks are structurally different: EU law permits balancing free expression against competing fundamental rights, while the U.S. First Amendment operates as a near-absolute constraint on government interference with speech. Because compliance with the AI Act is voluntary, however, the Supreme Court lacks the state action requirement needed to assess these constitutional effects, making judicial remedy unavailable. The constitutional tension therefore migrates to the legislative domain.

Key obligations of the AI Act would face serious First Amendment obstacles and would likely require substantial redesign before surviving constitutional review. But it produces comparable effects through market pressure anyway, and courts cannot remedy this because state action is absent, which means the pressure will ultimately fall on Congress.

Disclaimer: This thesis considers legislative developments until the 01.06.2026.

Abstract German:

KI-Systeme werden weltweit über zentralisierte Architekturen eingesetzt, die Nutzer gleichzeitig über verschiedene Rechtsordnungen hinweg bedienen. Das KI-Gesetz der Europäischen Union, der erste umfassende verbindliche Regulierungsrahmen für künstliche Intelligenz, reagiert darauf, indem er die Zuständigkeit für jeden Anbieter beansprucht, dessen Systeme Personen innerhalb der Union betreffen und das unabhängig vom Ort der Niederlassung. Das KI-Gesetz setzt dies durch den Binnenmarkt der EU um. Jeder Anbieter, der Nutzer innerhalb der Union bedienen möchte, muss die Vorschriften einhalten, unabhängig davon, wo dieser Anbieter ansässig ist. Die theoretischen Konsequenzen des sogenannten Brüssel-Effekts sind eindeutig: Unternehmen, einschließlich US-amerikanischer Firmen, entscheiden sich dafür, die EU-Compliance weltweit anzuwenden, anstatt ihre Betriebsabläufe zu fragmentieren. Dadurch kann die EU die globale KI-Governance allein durch Marktdruck über ihre Grenzen hinaus prägen.

Wenn US-Unternehmen die Verpflichtungen des EU KI-Gesetzes weltweit anwenden und dadurch EU-Regulierungen über Meinungsfreiheiten auf ihre gesamte Nutzerbasis übertragen, erzeugen diese Entscheidungen Auswirkungen auf die durch das „First Amendment“ geschützte Meinungsfreiheit, die nach der US-amerikanischen Verfassungsdoktrin eigentlich überprüft werden müssten. US-Gerichte wenden kein EU-Recht an. Sie überprüfen inländisches Verhalten, das durch ausländische Compliance-Anreize geprägt ist, und messen dieses an etablierten verfassungsrechtlichen Standards. Würde der Kongress vergleichbare Verpflichtungen verabschieden, würden dies politische und rechtliche Probleme im Hinblick auf das First Amendment schaffen. Die US-amerikanischen und europäischen Konzepte der Meinungsfreiheit unterscheiden sich strukturell. Während das EU-Recht eine Abwägung zwischen freier Meinungsäußerung und anderen Grundrechten zulässt, gilt das First Amendment in den USA nahezu als absolutes Verbot staatlicher Eingriffe in die Meinungsfreiheit. Da die Einhaltung des KI-Gesetzes jedoch freiwillig erfolgt, fehlt US-Gerichten das notwendige Merkmal des staatlichen Eingreifens „state action“, um diese verfassungsrechtlichen Auswirkungen zu prüfen, weshalb kein gerichtlicher Weg möglich ist. Die verfassungsrechtliche Spannung verlagert sich daher in den legislativen Bereich.

Mein zentrales Argument lautet, dass das KI-Gesetz verfassungswidrig sein könnte, wenn es vom US-Kongress verabschiedet würde, jedoch durch Marktdruck dennoch vergleichbare Wirkungen entfaltet. Gerichte können dagegen nichts unternehmen, weil das

erforderliche staatliche Handeln fehlt, weshalb der politische und rechtliche Druck letztlich auf den Kongress zurückfallen wird.

Hinweis: Diese Arbeit berücksichtigt gesetzgeberische Entwicklungen bis zum 01.06.2026.

Table of Contents:

I.	Introduction	1
II.	Conceptual foundations	3
III.	The EU AI Act as a source of transnational regulatory pressure	10
1.	The risk-based architecture	10
2.	Article 2 and the reach to third-country providers of EU regulation	12
3.	The global compliance mechanism and its dynamics	14
4.	Constitutional friction with U.S. values	15
IV.	AI Act's Obligations Under U.S. Constitutional Law	18
V.	How the AI Act shapes U.S. AI Governance in practice	22
1.	Market access pressure and the non-divisibility dynamic.....	22
a.	Why global standardisation is the rational commercial response	22
b.	The GDPR comparison: why AI Act compliance is architecturally deeper	24
2.	The compliance trajectory of U.S. firms	26
a.	Signals from early engagement: the GPAI Code of Practice.....	26
b.	The phased timeline and the compliance pressure still materialising	28
3.	Political resistance and its structural limits.....	30
VI.	The Failure of Judicial Remedies: Foreign Regulatory Diffusion and the Limits to Constitutional Litigation	33
1.	Why the Judicial Gap Is Structural and Not Incidental	33
a.	Voluntary Compliance and the Absence of Constitutional Review	33
b.	The Gap Between Constitutional Significance and Judicial Remedy	36
2.	The Conditions for Congressional Intervention.....	38
a.	The First Amendment Problem as a Leftover for Legislation	38
b.	Triggering Conditions and Available Regulatory Models.....	40
VII.	The Constitutional Gap and the Necessity of Congressional Action	44
VIII.	Conclusion	48

I. Introduction

What happens when the world's most ambitious artificial intelligence regulation collides with the world's most entrenched constitutional free speech culture? This thesis is an attempt to answer that question. The European Union's (EU) Artificial Intelligence (AI) Act, which entered into force in August 2024, represents the first comprehensive binding legal framework for AI technology. It is also, by design, a framework with global ambitions. Through its market access logic, the AI Act reaches any provider of AI systems whose outputs affect persons within the Union, regardless of where that provider is established. For firms operating in the United States (U.S.), this creates a practical compliance imperative that extends well beyond European borders. And it is precisely in that extension, the moment EU compliance obligations travel and shape the conduct of U.S. companies toward their domestic users, that something legally significant occurs.

A dominant theoretical account of this process is the Brussels Effect, which describes the EU's unilateral capacity to regulate global markets by shaping the compliance decisions of multinational firms. The logic rests on two conditions that both hold in the AI context: EU regulation is stricter than U.S. regulation, and firms can simultaneously comply with both. Where both conditions are met, the economics of compliance favour a single globally standardised product over parallel market-specific versions. It is important to be precise about what this means. Non-divisibility does not mean it is technically impossible to build separate AI models for the EU and U.S. markets. It means that producing and maintaining one globally compliant version is considerably cheaper than maintaining two, and so firms rationally extend EU standards worldwide rather than fragment their operations. The U.S., on this account, appears as largely neutral terrain across which EU regulation spreads without friction. This thesis challenges that assumption.

The U.S. is not a legally neutral surface. It is a constitutional order in which federal courts are empowered to assess the domestic legality of conduct regardless of its regulatory origin. When U.S. firms implement EU-derived content restrictions or transparency obligations across their entire user base, those decisions produce effects on First Amendment-protected expression that U.S. constitutional doctrine is built to evaluate. Courts do not apply EU law. What they review is domestic conduct that happens to have been shaped by foreign compliance incentives. The constitutional problem this creates is significant: the same obligations that the

AI Act lawfully imposes as conditions of EU market access would, if enacted by Congress, raise serious First Amendment difficulties. This is not incidental, as U.S. and EU free speech frameworks are structurally different. EU law permits free expression to be balanced against competing fundamental rights, while the U.S. First Amendment operates as a near-absolute constraint on government interference with speech. What is a permissible regulatory condition in one constitutional order is a potential constitutional violation in the other.

Yet the judicial remedy that this tension seems to call for is structurally unavailable. Because compliance with the AI Act is voluntary, meaning that firms choose to extend EU standards globally as a matter of commercial strategy and not because any U.S. government actor compels them to, the state action requirement that anchors First Amendment claims is systematically absent. Courts can only assess conduct where a legally cognizable injury exists and where government action can be traced. A private firm's decision to implement EU speech standards globally, however constitutionally significant its effects on U.S. users are, does not on its own satisfy that threshold. The constitutional tension is real, but it passes beneath the radar of judicial review. But what cannot be resolved through litigation does not disappear. It migrates. This thesis argues that the unresolved constitutional pressure generated by EU regulatory diffusion accumulates in the legislative domain, creating the conditions under which any durable domestic response would likely require congressional involvement.

My analysis proceeds in seven stages. First, I will establish the conceptual foundations (II): the Brussels Effect and the role of constitutional receiving orders in shaping how foreign regulatory norms take hold domestically. Then, this thesis will examine the EU AI Act as a source of transnational regulatory pressure, analysing its risk-based architecture, the market access logic of Article 2, and the constitutional friction that global compliance dynamics generate when they arrive in the U.S. context (III). What follows is a subsection of the AI Act's core obligations to direct U.S. constitutional analysis (IV). Afterwards, I will examine how the AI Act is already shaping U.S. AI governance in practice (V). Subsequently I will establish why judicial remedies are unavailable to address the constitutional tension that results (VI). After this, I will draw out the institutional implications, arguing for a resolution by domestic legislation (VII), followed by a conclusion (VIII).

II. Conceptual foundations

Understanding the transnational regulatory dynamics that this thesis analyses requires establishing three layers of conceptual groundwork: the distinctive character of AI as an object of regulation, the theoretical framework of the Brussels Effect, and the role of constitutional legal orders in shaping how foreign regulatory norms arrive and take hold domestically. These three layers intersect precisely at the point where EU AI regulation, diffused through private compliance decisions, produces effects on U.S. constitutional values that neither the Brussels Effect framework nor standard constitutional doctrines fully anticipate.

Digital technologies have long operated across jurisdictional boundaries. Online platforms, streaming services, and search engines serve users in multiple markets simultaneously through centralized architectures that do not, by design, observe national borders. In this sense, AI systems are not categorically different from other digital products. The technical infrastructure enabling global reach is broadly shared across the digital economy. The question worth pausing on is whether AI is transnational in a sense that is meaningfully distinctive, such that it generates regulatory problems of a qualitatively different character from those posed by earlier generations of global digital products. The answer does not lie in the architecture of delivery but in the nature of the risks AI systems produce. Unlike a streaming platform, which distributes content created by identifiable human authors, or a search engine, which indexes and ranks existing information, a generative AI system itself creates outputs across domains ranging from medical advice to legal analysis to political commentary, without necessarily any human editorial intervention in each instance. The system is trained on large datasets and generates content autonomously, making it both the instrument and, in a meaningful sense, the origin of the expressive conduct it produces.¹ This generative and autonomous character creates regulatory problems that earlier digital frameworks were not designed to address. The opacity of the processes producing the output, the difficulty of attributing responsibility for harmful content, and the potential for AI systems to cause harm at a scale and speed that human review cannot keep pace with.²

¹ Ryan Calo, 'Artificial Intelligence Policy: A Primer and Roadmap' (2017) 51(2) *UC Davis Law Review* <<https://digitalcommons.law.uw.edu/cgi/viewcontent.cgi?article=1639&context=faculty-articles>> p 405.

² Joanna J Bryson, Mihailis E Diamantis and Thomas D Grant, 'Liability for Artificial Intelligence' in Markus D Dubber, Frank Pasquale and Sunit Das (eds), *The Oxford Handbook of Ethics of AI* (Oxford University Press 20 October 2022) 187, <https://www.cambridge.org/core/services/aop-cambridge-core/content/view/12A89C1852919C7DBE9CE982B4DE54B7/9781009207867c12_187-209.pdf/liability-for-artificial-intelligence.pdf> p 190.

It is this combination of autonomy, opacity, scale, and fundamental rights implications that makes AI regulation a distinctive regulatory challenge, not merely the fact that AI systems cross borders. Many digital products do cross borders. What distinguishes AI is that the risks it produces are themselves difficult to contain jurisdictionally. An AI system trained on data from one jurisdiction and deployed through a centralized architecture to users across many jurisdictions simultaneously will produce outputs whose effects on speech, privacy, labour, education, and access to justice are felt wherever users are located. The tension between the territorial assumptions of legal jurisdiction and the borderless reality of digital harm is apparent.³ The EU's decision to regulate AI comprehensively through a binding horizontal framework responds to this risk character, not simply to the fact of global reach.⁴

The dominant theoretical account of how EU regulation shapes the conduct of firms operating outside EU territory is the Brussels Effect, defined by Anu Bradford as the EU's unilateral ability to regulate the global marketplace through the compliance decisions of multinational firms. The concept describes a process in which firms operating in the EU market, finding that EU standards are stricter than those of other jurisdictions, choose to extend EU compliance globally rather than maintain separate product versions for separate markets.⁵ The result is that EU regulatory standards acquire a global footprint not through treaty obligation or enforcement across borders, but through the private rational choices of firms seeking market access. It is important to establish at the outset that the Brussels Effect is not a distinctively European phenomenon in its underlying logic. The mechanism depends on market size.⁶ Any sufficiently large market whose regulator sets standards stricter than those prevailing elsewhere will generate comparable dynamics, provided firms cannot easily serve that market while ignoring its requirements.⁷ The EU generates the Brussels Effect in practice because it constitutes the world's largest single market with credible regulatory enforcement capacity, and because major technology firms cannot commercially afford to forgo access to it.⁸ The

³ Jennifer Daskal, 'Borders and Bits' (January 2018) 71(1) *Vanderbilt Law Review* 179, <<https://www.proquest.com/docview/2040740630/fulltextPDF/D356EA0C515473BPQ/1?accountid=14682&source=Scholarly%20Journals>> p 186.

⁴ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, p 1-2.

⁵ Anu Bradford, 'Part front matter for Preface to Part One Theory' in *The Brussels Effect: How the European Union Rules the World* (1st edn, Oxford University Press 2020) p 2.

⁶ Chad Damro, 'Market Power Europe' (1 March 2012) 19(5) *Journal of European Public Policy* 682, <<https://doi.org/10.1080/13501763.2011.646779>> p 690.

⁷ Anu Bradford, 'Globalizing European Digital Rights through Regulatory Power' in *Digital Empires: The Global Battle to Regulate Technology* (New York, 2023; online edn, Oxford Academic, 21 September 2023) p 328.

⁸ Anu Bradford, 'The Brussels Effect' (2012) 107(1) *Northwestern University Law Review* <<https://www.proquest.com/docview/1286682884?parentSessionId=h1I5qUw%2F5GpZbiktLEx6LUDZu%2FD>

theoretical mechanism would operate identically if the U.S. were to enact comparably strict AI regulation. U.S. standards would then travel globally through the same market access logic. Two conditions must hold simultaneously for the Brussels Effect to operate in any given regulatory context. First, the EU regulatory standard must be stricter than that prevailing in other jurisdictions, including the U.S., such that compliance with EU law goes beyond what domestic law requires. Second, a firm must be able to comply with both EU and domestic law simultaneously. This means that EU obligations must not require conduct that U.S. law prohibits, or vice versa.⁹ Where both conditions hold, global standardisation around the EU standard becomes the rational commercial choice. Where either condition breaks down, the dynamic changes fundamentally. If EU and U.S. standards converge, the incentive to globalise EU compliance disappears. If complying with the AI Act required violating U.S. law, firms would face a genuine conflict rather than a compliance choice, and the Brussels Effect would be blocked.

In the AI context these conditions seem to hold up, though not without tension. The AI Act imposes obligations on the design, documentation, transparency, and human oversight of AI systems that go significantly beyond anything currently required under U.S. federal law.¹⁰ Firms can, at least in principle, comply with the AI Act without violating any U.S. legal obligation, because U.S. law does not prohibit the conduct the AI Act requires at present. The constitutional tension this thesis analyses arises precisely in the gap between legal permissibility and constitutional desirability. Compliance is voluntary and lawful, but the obligations complied with, if imposed by Congress in comparable ways with the AI Act, would raise serious constitutional questions.

Central to the Brussels Effect is the concept of non-divisibility, which is the observation that firms often find it more efficient to operate a single globally uniform product than to maintain separate versions tailored to the requirements of different markets.¹¹ This concept requires careful statement. Non-divisibility does not mean it is technically impossible to build separate AI models for different markets. Firms could produce one AI system for the EU market, calibrated to AI Act requirements, and a different system for the U.S. market, governed

[dXTJRjDyMxCgK9uk%3D&pq-origsite=primo&accountid=14682&sourcetype=Scholarly%20Journals](#)> p 25-28.

⁹ Charlotte Siegmann and Markus Anderljung, ‘The Brussels Effect and Artificial Intelligence: How EU Regulation Will Impact the Global AI Market’ (Governance of AI Research Paper, 16 August 2022) <https://cdn.governance.ai/Brussels_Effect_GovAI.pdf> p 27-28.

¹⁰ *Ibid*, 33-34.

¹¹ Bastiaan Crum, ‘Brussels Effect or Experimentalism? The EU AI Act and Global Standard-Setting’ (27 August 2025) 14(3) *Internet Policy Review* <<https://policyreview.info/pdf/policyreview-2025-3-2032.pdf>> p 3.

only by domestic U.S. standards. The technical capacity to do this exists, even if it is economically irrational under conditions of high compliance cost and strong market access incentives.¹² Non-divisibility describes an economic constraint and not a technical one. Maintaining two parallel product architectures, each with its own training pipeline, compliance documentation, testing regime, and infrastructure, is considerably more expensive than maintaining one. Where the less demanding standard can be satisfied by the product designed to meet the more demanding one, the rational commercial choice is to produce a single globally compliant version and avoid the cost of duplication.¹³ This distinction matters for the analysis that follows. Because the decision to globalise EU compliance is an economic choice made by the firm, not a technical necessity, the firm is a genuine agent in the process. It chooses to extend EU standards to U.S. users. No U.S. government actor compels it to do so. This voluntary character of the compliance decision is precisely why the state action requirement that would anchor constitutional claims in the U.S is systematically absent from the scenarios analyzed, and why courts are structurally limited in their capacity to assess the constitutional effects that result in their own jurisdiction.

The Brussels Effect operates through two analytically distinct mechanisms, each of which produces a different relationship between EU regulatory standards and the legal order of the jurisdiction in which they arrive. The de jure Brussels Effect describes the process by which foreign governments actively decide to adopt EU regulatory standards into their own legal frameworks, sometimes in response to domestic lobbying by firms that prefer a single global standard over the cost of managing regulatory divergence.¹⁴ In the de jure case, EU standards become domestic law through a sovereign legislative act, and so the receiving jurisdiction chooses to replicate EU requirements. The General Data Protection Regulation (GDPR) has produced de jure effects in jurisdictions including Brazil¹⁵, Japan¹⁶, and South Korea¹⁷, where data protection legislation has been inspired by or modelled on EU requirements. The de facto Brussels Effect is the more analytically significant mechanism for this thesis. It describes the process by which EU rules acquire global force not because governments actively choose to

¹² Anu Bradford, 'Globalizing European Digital Rights through Regulatory Power' in *Digital Empires: The Global Battle to Regulate Technology* (New York, 2023; online edn, Oxford Academic, 21 September 2023) p 327.

¹³ fn (11).

¹⁴ Anu Bradford, 'The Brussels Effect in Context' in *The Brussels Effect: How the European Union Rules the World* (1st edn, Oxford University Press 2020) p 78.

¹⁵ Paul M Schwartz, 'Global Data Privacy: The EU Way' (2019) 94(4) *New York University Law Review* 771 <<https://nyulawreview.org/wp-content/uploads/2019/10/NYULAWREVIEW-94-4-Schwartz.pdf>> p 812.

¹⁶ *Ibid*, p 789.

¹⁷ *Ibid*, p 807.

adopt them, but because firms voluntarily extend compliance beyond EU borders to maintain market access.¹⁸ A firm applies EU standards to its entire global operation because doing so is cheaper than segmenting by jurisdiction. The regulatory diffusion occurs through private corporate decisions, without any legislative act in the receiving jurisdiction, and without any formal extension of EU legal authority beyond EU territory. It is the de facto mechanism that produces the constitutional dynamics this thesis examines. When U.S. AI providers extend AI Act compliance globally, they do so through voluntary commercial decisions. Those decisions produce effects on First Amendment protected expression in the U.S that are legally significant, but that arise from private conduct rather than government mandate. The constitutional question this generates is there, but its resolution through litigation is structurally blocked, for reasons that the following chapters develop in detail. It is this structural feature of voluntary diffusion, rather than any weakness in the constitutional arguments themselves, that makes the judicial route unavailable and shifts the burden to the legislative domain.

The Brussels Effect is not the only theoretical framework through which the AI Act's global reach can be understood. Situating it within a broader view strengthens rather than weakens the analysis. For example, experimentalist governance argues that the AI Act's relatively open textured and procedural character invites an iterative, learning based regulatory dynamic that differs from the classical Brussels Effect model of standard setting followed by global adoption.¹⁹ On this account, the AI Act does not simply export fixed EU standards that third country firms adopt wholesale. Instead, it creates a framework within which regulatory norms are developed iteratively through implementation, enforcement, and dialogue between regulators, firms, and civil society, with the possibility that the standards themselves evolve in response to experience.²⁰ The implication for global diffusion is that what travels through the Brussels Effect mechanism may be not a fixed regulatory template but a dynamic regulatory process, one that is more adaptable to local constitutional conditions than the classical model suggests. The experimentalist frame is useful as it reinforces the point that the AI Act's global influence is conditional rather than automatic. Diffusion depends on the absence of structural barriers in receiving jurisdictions, and the presence of constitutional resistance in the U.S. is

¹⁸ Anu Bradford, 'The Brussels Effect' (2012) 107(1) *Northwestern University Law Review* <<https://www.proquest.com/docview/1286682884?parentSessionId=hI15qUw%2F5GpZbiktLEx6LUDZu%2FDdXTJRjDyMxCgK9uk%3D&pq-origsite=primo&accountid=14682&sourcetype=Scholarly%20Journals>> p 7.

¹⁹ Bastiaan Crum, 'Brussels Effect or Experimentalism? The EU AI Act and Global Standard-Setting' (27 August 2025) 14(3) *Internet Policy Review* <<https://policyreview.info/pdf/policyreview-2025-3-2032.pdf>> p 5.

²⁰ *Ibid*, p 16.

precisely one such barrier,²¹ even if not a legal one as of now. Also, it cautions against treating the Brussels Effect as a clear outcome of either full global adoption or complete failure. The more realistic picture, which this thesis's analysis of the U.S. case supports, is one of selective, partial, and constitutionally mediated diffusion, in which some AI Act obligations travel through compliance decisions while others encounter barriers that redirect the dynamic toward legislative rather than market-based resolution. The Brussels effect is still used as the reference in this thesis, but it is inherently only one type of view of regulatory diffusion and not an absolute.

The Brussels Effect framework might seem to treat receiving jurisdictions largely as passive actors across which EU regulatory norms spread through market pressure, or as secondary actors. This is theoretically incomplete. Jurisdictions are not legally neutral surfaces, but constitutional orders equipped with their own doctrinal tools for assessing the legality of conduct that occurs within their territory, regardless of the regulatory origin of that conduct.²² In the U.S, this means that when firms apply AI Act compliance obligations to their domestic users, the resulting conduct is subject to assessment under the U.S. Constitution, applied by U.S. courts, according to U.S. constitutional doctrine. Courts do not apply EU law or review the AI Act. What courts review is domestic conduct that happens to have been shaped by a foreign compliance incentive, measured against settled domestic constitutional standards. This dynamic has already been seen before, as with the GDPR. U.S. courts there would assess the conduct of U.S. firms that adopted GDPR standards applying U.S. law to U.S. conduct but not EU law. The foreign regulatory origin of the conduct was analytically relevant but not as a source of legal obligation. The constitutional doctrine most directly implicated by AI Act compliance is the First Amendment doctrine. The AI Act imposes content restrictions and compelled disclosure obligations that, if enacted by Congress, would face significant scrutiny under the First Amendment.²³ Other constitutional questions focused on due process for example also come into question, but this thesis will limit itself to First Amendment concerns. More precisely, U.S. and EU free speech frameworks differ at a foundational level. EU law treats freedom of expression as one fundamental right among others, subject to balancing against competing rights and public interests.²⁴ The U.S. First Amendment operates as a near

²¹ Marco Almada and Anca Radu, 'The Brussels Side-Effect: How the AI Act Can Reduce the Global Reach of EU Policy' (2024) 25 German Law Journal 657, p 661.

²² Anu Bradford, 'The Brussels Effect' in *The Brussels Effect: How the European Union Rules the World* (1st edn, Oxford University Press 2020) p 31.

²³ US Constitution First Amendment.

²⁴ Charter of Fundamental Rights of the European Union [2012] OJ C 326/391, arts 11 and 52(1).

absolute constraint on government interference with expression, one that resists the kind of balancing that EU constitutional law readily permits.²⁵ What is a lawful regulatory condition under EU constitutional law may be an unconstitutional government imposition under U.S. doctrine.

The critical point for this chapter's argument is that this constitutional incompatibility does not translate automatically into judicial review. The Supreme Court can only assess conduct where a legally cognizable injury exists and where state action can be established.²⁶ A private firm's decision to extend EU compliance to its U.S. operations, however constitutionally significant its effects on U.S. users, does not satisfy the state action requirement that anchors constitutional claims under the First Amendment. The constitutional tension is real as the judicial remedy is structurally unavailable. And what cannot be resolved through litigation accumulates as pressure in the legislative domain, creating the conditions under which Congress will need to act. It is this combination of constitutional significance, judicial unavailability, and legislative necessity that defines the analytical contribution of the chapters that follow. The Brussels Effect does not operate seamlessly in the U.S. It encounters a constitutional legal order that is structurally positioned to shape the terms on which foreign regulatory norms arrive and take hold, and it finds that the mechanism for doing so, being judicial review, is systematically absent in the voluntary compliance context. What results is not the frictionless global diffusion the Brussels Effect model might anticipate, but a conditional, institutionally mediated process whose endpoint is legislative rather than judicial.

²⁵ Frederick Schauer, 'The Exceptional First Amendment' (KSG Working Paper No RWP05-021, 18 February 2005) <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=668543> p 4.

²⁶ Priscilla H Machado, 'State Action and the Supreme Court' (EBSCO Research Starters, 2023) <<https://www.ebsco.com/research-starters/law/state-action-and-supreme-court>>.

III. The EU AI Act as a source of transnational regulatory pressure

Having established the conceptual foundations of the Brussels Effect and the constitutional distinctiveness of the U.S. as a receiving jurisdiction, this chapter turns to the EU AI Act itself. The aim is threefold. First, to analyse the internal legal architecture through which the Act imposes obligations, second to identify the market logic through which it reaches providers established outside the Union, and third to examine the global compliance dynamics that follow, including the constitutional friction those dynamics produce in the U.S. context. The AI Act does not pressure third countries through explicit extraterritorial mandates. It does so structurally, by making access to the EU market conditional on compliance with a deeply architecturally demanding framework that, by its internal logic, cannot easily be confined to EU operations alone.

1. The risk-based architecture

The AI Act organises its regulatory intervention around a tiered classification of AI systems according to the level of risk they pose to safety, fundamental rights, and public interests. This classification is not the product of a single provision. It emerges from the combined operation of several articles and annexes read together.²⁷ The strictest classification is constituted by a set of practices prohibited outright under Article 5 of the AI Act, which lists eight categories of AI use deemed incompatible with EU values and therefore impermissible²⁸ regardless of the provider's origin or the intended purpose according to Article 2.²⁹ These include AI systems that deploy subliminal techniques to manipulate behaviour, exploit vulnerabilities of specific groups, conduct real-time remote biometric identification in publicly accessible spaces subject to limited exceptions, and create or expand facial recognition databases through untargeted scraping. The prohibitions on these practices became enforceable in February 2025, pursuant to Article 113(a) of the AI Act, which provides for their application

²⁷ European Commission, 'Regulatory Framework Proposal on Artificial Intelligence' <<https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>>.

²⁸ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, art 5.

²⁹ *Ibid*, art 2.

six months after entry into force.³⁰ Below the prohibited systems, the Act identifies high-risk AI systems in Articles 6(1) and 6(2). Article 6(1) covers AI systems that are safety components of products subject to Union harmonisation legislation listed in Annex I.³¹ Article 6(2) covers AI systems listed in Annex III, which enumerates eight domains including biometric identification, critical infrastructure, education, employment, access to essential services, law enforcement, migration management, and the administration of justice.³²

It is this classification that triggers the substantive compliance obligations laid out in Articles 9 through 15 of the Act. These provisions require providers of high-risk systems to establish a risk management system under Article 9, to implement data governance practices under Article 10, to produce technical documentation under Article 11, to record-keep under Article 12, to ensure transparency toward deployers under Article 13, to enable meaningful human oversight under Article 14, and to meet accuracy, robustness and cybersecurity standards under Article 15.³³ A critical clarification must be made here. Articles 9 through 15 apply to all high-risk AI systems placed on or put into service in the EU, whether developed by EU or third-country providers. They are not directed specifically at foreign AI systems. Their disproportionate practical impact on U.S. providers is a consequence of the current composition of the global AI market, in which the dominant general-purpose AI systems, large language models, and high-capability frontier models are predominantly developed by U.S. companies, such as Anthropic or OpenAI, owning respectively the AI systems Claude and ChatGPT.³⁴ The obligations fall on all providers equally and so the market reality is that many providers subject to those obligations are not European, which puts the focus on the impact of the AI Act on foreign providers.³⁵

The Act also establishes a limited transparency tier under Article 50, which requires providers of certain AI systems to ensure that users are informed they are interacting with an AI, and that AI-generated content is marked as such.³⁶ This tier applies more broadly than the

³⁰ *Ibid*, art 113(a).

³¹ *Ibid*, art 6(1).

³² *Ibid*, Annex III.

³³ *Ibid*, arts 9 and 15.

³⁴ Anthropic PBC is incorporated in Delaware and headquartered in San Francisco, California; OpenAI is incorporated in Delaware and headquartered in San Francisco, California. Both are established US companies cited for the purposes of this analysis.

³⁵ Charlotte Siegmann and Markus Anderljung, ‘The Brussels Effect and Artificial Intelligence: How EU Regulation Will Impact the Global AI Market’ (Governance of AI Research Paper, 16 August 2022) <https://cdn.governance.ai/Brussels_Effect_GovAI.pdf> p 33-34.

³⁶ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, art 50.

high-risk category and imposes lighter obligations. With respect to minimal risk systems, the AI Act does not impose mandatory compliance obligations. The absence of mandatory obligations for this category does not mean the category is unaddressed. It means the Act affirmatively declines to regulate it through binding rules. This is relevant for understanding the Act's scope. The majority of AI systems deployed in the EU today fall below the high-risk threshold and are therefore not directly regulated by the Act's substantive provisions.

The risk-based architecture matters in connection with the Brussels Effect because it defines what compliance entails. Compliance with Article 5 prohibitions requires firms to refrain from specific deployment practices. Compliance with Articles 9 through 15 requires something far more demanding, being the structural redesign of training systems, documentation systems, testing procedures, and deployment architectures. It is this depth of required compliance that distinguishes the AI Act from prior EU regulatory frameworks and that makes the non-divisibility logic particularly consequential.

2. Article 2 and the reach to third-country providers of EU regulation

If the risk-based architecture defines what the AI Act requires, Article 2 defines who it requires it of. The Act's market logic is established by Article 2(1), stating the following³⁷:

*“This Regulation applies to: (a) providers placing on the market or putting into service AI systems or placing on the market general-purpose AI models in the Union, **irrespective of whether those providers are established or located within the Union or in a third country**;*

*(c) providers and deployers of AI systems that have **their place of establishment or are located in a third country**, where the **output produced by the AI system is used in the Union**;*

*(f) authorised **representatives of providers, which are not established in the Union**;*

*(g) affected **persons that are located in the Union**.*”

The term "provider" is defined as a natural or legal person, public authority, agency or other body that develops an AI system or a general-purpose AI model and places it on the

³⁷ *Ibid*, art 2(1).

market or puts it into service under its own name or trademark, whether for payment or free of charge.³⁸ This definition captures U.S. firms that develop and deploy AI systems accessed by persons within the Union, including through online interfaces, regardless of the server location or corporate establishment of the developer. The mechanism through which Article 2 generates compliance pressure for third-country providers operates through market access conditionality. A U.S. AI provider is under no obligation under EU law to serve the European market. But if it wishes to and does, full compliance with the Act's applicable provisions is legally required, and non-compliance exposes the provider to the enforcement and sanction regime established by the Act. The primary compliance mechanism is legal rather than voluntary. Article 99 of the AI Act provides for fines of up to EUR 35 million.³⁹ This enforcement framework applies to all providers subject to the Act, including those established in third countries.

To facilitate enforcement against providers not established in the Union, Article 22 of the Act provides that any provider not established in the Union shall, before placing an AI system on the Union market or putting it into service in the Union, designate an authorised representative established in the Union.⁴⁰ The representative must be able to perform the tasks specified in the mandate, and the provider remains jointly responsible for compliance. This is a mandatory legal obligation, not a voluntary option. It creates a formal legal connection between the third-country provider and Union regulatory authority, ensuring that enforcement is not weakened by the provider's geographic location, but also intervening structurally as it demands the firms to have an institutional foothold within EU territory as a condition of market access.

The GDPR offers a useful comparative reference for understanding how this market logic operates in practice and demonstrates that the AI Act's foreign reach is not unprecedented. It similarly extends its application to controllers and processors not established in the Union where they offer goods or services to data subjects in the Union or monitor their behaviour.⁴¹ The position is more complex than simple exemption or inclusion, because U.S. firms that process EU personal data are subject to the GDPR and must comply with its transfer mechanisms and substantive rules. The EU-U.S. Data Privacy Framework established in 2023 creates a compliance pathway for certified U.S. organisations, but it does not exempt them from

³⁸ *Ibid*, art 3.

³⁹ *Ibid*, art 99(3).

⁴⁰ *Ibid*, art 22.

⁴¹ 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) [2016] OJ L 119/1, art 3(2).

GDPR requirements. It provides a mechanism for demonstrating compliance with the conditions for international data transfers.⁴² The AI Act's reach is analogous in structure but architecturally deeper in consequence. Where GDPR compliance required firms to adjust data handling practices, AI Act compliance for high-risk systems under Articles 9 through 15 requires restructuring the entire design, documentation, testing and ongoing monitoring of entire AI systems. This matters because it shapes the intensity of the Brussels Effect. Firms that comply with high-risk obligations do not add a compliance layer to an existing product but they build compliance into the product's architecture.

3. The global compliance mechanism and its dynamics

The enforcement-backed compliance framework of Articles 2 and 22, combined with the depth of obligations under Articles 9 through 15, produces a global compliance dynamic that extends beyond the EU's territory through the mechanism of non-divisibility discussed in Chapter II. It is here that the Act begins to function as a source of transnational pressure, not by commanding behaviour in third countries through an extraterritorial effect, but by creating strong structural incentives for firms to extend EU compliance globally rather than maintain separate operational architectures. Article 113 of the AI Act establishes the phased timeline for these obligations, dividing the entrance into force of the Act into multiple phases. Article 5 prohibition provisions apply from February 2025, obligations relating to general-purpose AI models under Chapter V apply from August 2025 and the remaining substantive provisions, including the high-risk AI obligations of Articles 9 through 15, apply from August 2026. Certain provisions relating to AI systems that are safety components even apply only from August 2027.⁴³ This phased timeline is significant because it means the full architecturally demanding compliance pressure is only beginning to materialise for most U.S. providers. The General-Purpose AI Code of Practice (GPAI), published in July 2025, illustrates the emerging dynamic. The Code is not legally mandatory but addresses transparency, copyright, and safety obligations for general-purpose AI model providers under Chapter V of the Act.⁴⁴ Several major

⁴² US Department of Commerce, 'Data Privacy Framework Program Overview' <<https://www.dataprivacyframework.gov/Program-Overview>>.

⁴³ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, art 113.

⁴⁴ European Commission, 'The General-Purpose AI Code of Practice' <<https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>>.

U.S. AI providers, including Anthropic and OpenAI, are among its signatories.⁴⁵ Signature does not constitute compliance with the full AI Act, and in particular does not satisfy the high-risk obligations of Articles 9 through 15. However, it signals that leading U.S. providers are engaging with EU regulatory standards at the design and governance level, which is consistent with the de facto Brussels Effect mechanism described in Chapter II.

The Trump administration's executive order of January 2025, revoking the Biden AI executive order of October 2023⁴⁶, introduces a complicating political dimension. The revocation does not alter the legal obligation of U.S. firms to comply with the AI Act if they wish to operate in the EU market. That obligation derives from EU law and is independent of U.S. domestic policy. What the executive order does affect is the political and institutional environment in which compliance decisions are made. A federal government that actively discourages AI regulation and signals that foreign regulatory standards conflict with national policy⁴⁷ reduces the institutional incentives for firms to treat EU compliance as a model for domestic operations. The more significant consequence concerns the systemic resistance observed. If U.S. regulatory policy creates a sustained divergence from EU standards, firms operating in both markets may ultimately face a choice between maintaining two architecturally distinct AI systems, one EU-compliant and one not, or withdrawing from one market. This potential fragmentation is precisely what the non-divisibility logic of the Brussels Effect argues against, and its emergence would represent a genuine structural limit on the Act's capacity to generate global standardisation.

4. Constitutional friction with U.S. values

The compliance dynamics described previously do not arrive in the U.S. as legally neutral phenomena. When U.S. AI providers extend AI Act obligations to their domestic operations through unified compliance architectures, specific categories of obligation engage specific areas of U.S. constitutional sensitivity. It is important to be precise about which obligations produce this friction and why, because the constitutional concern is not a general

⁴⁵ *Ibid.*

⁴⁶ Executive Office of the President, 'Removing Barriers to American Leadership in Artificial Intelligence' (White House, 23 January 2025) <<https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/>>.

⁴⁷ *Ibid.*, section 1.

unease about foreign regulatory influence. It is a doctrinal problem that arises from the precise content of specific AI Act provisions when those provisions operate in a legal environment built around the First Amendment.

Two categories of AI Act obligation are constitutionally significant in the U.S. context, and they raise distinct problems. Each raises distinct doctrinal problems, which Chapter IV subjects to direct constitutional analysis. The first category consists of the content restrictions in Article 5.⁴⁸ The eight practices prohibited by Article 5 are prohibited because of what they do. They produce, deploy, or facilitate AI outputs that the EU legislature determined are incompatible with fundamental rights and human dignity.⁴⁹ The prohibitions target the content and purpose of AI outputs directly. When a U.S. provider extends these prohibitions to its domestic user base as part of a unified compliance strategy, the result is that U.S. users interact with AI systems whose output categories have been defined by EU value judgments about permissible expression. Those judgments were calibrated to the EU fundamental rights framework and the proportionality standard it applies. They were not designed to satisfy the strict scrutiny that the First Amendment imposes on content-based restrictions.⁵⁰ A U.S. AI provider operating under Article 5 compliance standards is making editorial choices about what its system will and will not generate which then reflect EU constitutional choices rather than U.S. ones. The provider is entitled to make those choices, but what is constitutionally significant is that the standards shaping them originate in a foreign legislature rather than in U.S. democratic deliberation, and no U.S. constitutional test was applied to them or, as we will analyse further, can be.

The second category consists of the transparency and disclosure obligations in Articles 13 and 50. Article 50 requires providers to inform users that they are interacting with an AI system and to ensure that AI-generated content is identified as such.⁵¹ Article 13 requires providers of high-risk systems to supply deployers with information sufficient to understand the system's outputs, performance characteristics, and limitations.⁵² These obligations do not restrict what AI systems may generate. They require AI providers to produce specific disclosures about how their systems operate. In U.S. constitutional terms, this is compelled

⁴⁸ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, art 5.

⁴⁹ Marco Almada and Nicolas Petit, 'The EU AI Act: A Medley of Product Safety and Fundamental Rights?' (Robert Schuman Centre for Advanced Studies Research Paper No 2023/59, 18 October 2023) <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4308072> p 8.

⁵⁰ US Constitution First Amendment.

⁵¹ *fn* (48), art 50 (1).

⁵² *Ibid*, art 13.

speech.⁵³ The government can require companies to provide factual and noncontroversial information in commercial settings, which will be developed later in this thesis. However, when companies are forced to explain how AI systems work, especially when those systems involve disputed technical or ethical questions, there must be a stronger constitutional justification for requiring that disclosure. Article 13's requirement that providers explain the limitations and accuracy characteristics of high-risk AI systems⁵⁴ goes beyond a simple factual notice. It requires providers to produce interpretive characterisations of systems that are not fully transparent probably even to their own developers, and to do so in terms that reflect the EU's risk classification framework and the value judgments embedded in it. If these requirements were enacted as U.S. legislation, they would face challenge on the ground that they compel speech that goes beyond what the permissive standard for commercial disclosure requirements allows.

Taken together, these two categories of friction share a common structural feature. The obligations in question are not unreasonable within the EU's constitutional order. They rest on a framework that takes fundamental rights seriously and applies proportionality analysis carefully. The problem is that when those obligations migrate into U.S. operational practice through voluntary compliance, they bring the EU constitutional framework with them, applied to U.S. users, without any of the constitutional scrutiny that U.S. law would require if the same obligations were imposed nationally. The content choices of AI systems operating in the U.S. market are being shaped by EU speech values. The disclosure requirements those systems produce for U.S. users are calibrated to EU informational standards. Neither was subjected to the First Amendment analysis that U.S. constitutional culture appears to demand.

⁵³ Congress of the United States of America, 'Amdt1.7.14.1 Overview of Compelled Speech' in *Constitution Annotated* <https://constitution.congress.gov/browse/essay/amdt1-7-12-1/ALDE_00000769/>.

⁵⁴ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, art 13(3)(b).

IV. AI Act's Obligations Under U.S. Constitutional Law

Having established the AI Act's obligations and the market logic through which it reaches U.S. providers, this chapter considers if an equivalent set of obligations were enacted by Congress, how U.S. constitutional doctrine would assess them. The answer matters not because Congress has enacted such a statute, as it has not done that at present, but because the obligations U.S. firms voluntarily extend to their domestic operations through unified compliance architectures produce effects on First Amendment-protected expression that are structurally similar to those a domestic statute would produce. The constitutional analysis therefore serves two purposes simultaneously. It establishes what U.S. constitutional law would require of equivalent domestic legislation, and it maps the gap between those requirements and the standards applied when private firms absorb EU obligations through voluntary compliance.

Before applying First Amendment doctrine to the specific categories of AI Act obligation, it is necessary to establish the structural divergence between U.S. and EU free speech frameworks, because that divergence is not incidental. It is the precondition for the constitutional problem this thesis examines. Freedom of expression is guaranteed by both the U.S. Constitution and the EU Charter of Fundamental Rights, but the two frameworks differ fundamentally in their relationship to competing rights and public interests.⁵⁵ Under Article 11 of the Charter, freedom of expression is a fundamental right, but Article 52(1) expressly permits limitations on Charter rights where those limitations are “provided for by law”, respect the essence of the rights, are necessary, and “genuinely meet objectives of general interest recognized by the Union or the need to protect the rights and freedom of others.”⁵⁶ The result is a proportionality-based framework in which free expression is regularly and legitimately balanced against competing values such as human dignity, non-discrimination, public health, and the protection of democratic processes. The AI Act's content prohibitions in Article 5 are products of this balancing. The EU legislature determined that certain AI-enabled practices are incompatible with fundamental rights and restricted them accordingly after proportionality analysis calibrated to the EU constitutional order.⁵⁷

⁵⁵ Frederick Schauer, ‘The Exceptional First Amendment’ (KSG Working Paper No RWP05-021, 18 February 2005) <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=668543> p 33-34.

⁵⁶ Charter of Fundamental Rights of the European Union [2012] OJ C 326/391, arts 11 and 52(1).

⁵⁷ fn (54), art 5.

The U.S. First Amendment operates on a different structural basis. The Supreme Court has consistently refused to adopt a general balancing approach to speech restrictions, holding instead that content-based restrictions on speech imposed by the government are presumptively unconstitutional and subject to strict scrutiny.⁵⁸ In *Reed v. Town of Gilbert* in 2015, the Court confirmed that even if a law does not target a specific opinion, it is still considered content based if it blocks public discussion of an entire topic. Under the First Amendment, such a law must meet the highest level of constitutional review, called strict scrutiny.⁵⁹ Although the *Reed* doctrine is criticised and sometimes labeled as too strict⁶⁰, there is no proportionality stage in U.S. First Amendment doctrine comparable to Article 52(1) of the Charter. The government cannot justify a content restriction by demonstrating that competing interests outweigh free expression, but it must demonstrate that the restriction is the least restrictive means of serving an interest.⁶¹ This structural divergence explains why obligations that are constitutionally unproblematic within the EU regulatory order may be constitutionally problematic if enacted domestically in the U.S. What the AI Act achieves through proportionality, U.S. legislation would need to achieve through strict scrutiny, which is a considerably more demanding threshold.

Furthermore, Article 5 of the AI Act prohibits eight categories of AI practice outright, including systems that deploy subliminal or manipulative techniques to distort behavior, systems that exploit the vulnerabilities of specific groups, and systems that enable real-time remote biometric identification in publicly accessible spaces.⁶² These prohibitions are defined by reference to the purpose and content of AI outputs, meaning what the system does and what it produces⁶³, rather than by reference to the manner or time of production. Under settled First Amendment doctrine, a restriction defined by reference to the content or communicative purpose of expression is content-based, regardless of whether it is framed as a safety measure linked to health-care.⁶⁴ If Congress were to enact Article 5's prohibitions as domestic national legislation, they would face two immediate First Amendment difficulties. The first concerns the category of “subliminal” manipulation. Article 5(1)(a) prohibits AI systems that “deploy

⁵⁸ *Grayned v City of Rockford* 408 US 104 (1972).

⁵⁹ *Reed v Town of Gilbert* 576 US 155 (2015).

⁶⁰ Alexander Tsesis, ‘Compelled Speech and Proportionality’ (2022) 97(3) *Indiana Law Journal* art 1 <<https://www.repository.law.indiana.edu/ilj/vol97/iss3/1>> p 819.

⁶¹ *United States v Playboy Entertainment Group Inc* 529 US 803 (2000).

⁶² Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, art 5.

⁶³ *Ibid*, art 5(1)(h)(iii).

⁶⁴ *Sorrell v IMS Health Inc* 564 US 552 (2011).

subliminal techniques beyond a person's consciousness” with the objective of materially distorting that person's behaviour in a manner that causes or is likely to cause harm.⁶⁵ A U.S. statute restricting AI outputs on the ground that they are manipulative or distorting would be subject to strict scrutiny being a content-based restriction targeting expressive conduct based on its perceived effect on the audience. The Supreme Court has rejected the argument that speech may be restricted because its persuasive or affective power is harmful. For example, in *Snyder v. Phelps*, the Court protected offensive expressive conduct because content-based liability for speech that causes emotional harm cannot survive First Amendment scrutiny.⁶⁶ The second difficulty concerns the facial breadth of the prohibitions. Several of Article 5's categories depend on concepts such as 'manipulative techniques,' 'vulnerable groups,' 'significant harm', whose application most likely requires a case-by-case judgment that in U.S. doctrine creates risks of vagueness if also enacted in the same way. A law that bans speech based on vague or subjective standards may be unconstitutional, even if some applications are valid, because it can discourage more speech than the government is allowed to regulate.⁶⁷ The point is not that the practices Article 5 prohibits are desirable in the U.S. or that the EU's judgment about their incompatibility with fundamental rights is wrong. The point is structural in the sense that the First Amendment does not permit the U.S. government to reach those prohibitions through the same doctrinal pathway the EU employed. A domestic equivalent would require either a different legal architecture or constitutional justification it may not be possible to supply on the same grounds.

The AI Act's transparency rules in Articles 13 and 50 raise related First Amendment concerns. Article 50 requires users to be told when they interact with AI or AI-generated content⁶⁸, while Article 13 requires providers of high-risk systems to disclose information about the system's capabilities, accuracy, and limitations.⁶⁹ These provisions do not restrict what AI systems may generate. They compel AI providers to produce specific statements about their systems and the content those systems generate. In First Amendment terms, they are instances of compelled speech. The constitutional framework for compelled speech in commercial contexts, which would also apply to the consumer character of AI system users, originates in *Zauderer v. Office of Disciplinary Counsel*, which held that the government may require

⁶⁵ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, art 5(1)(a).

⁶⁶ *Snyder v. Phelps* 562 US 443 (2011).

⁶⁷ Congress of the United States of America, 'Amdt1.7.2.2 Vagueness, Statutory Language, and Free Speech' in *Constitution Annotated* <https://constitution.congress.gov/browse/essay/amdt1-7-2-2/ALDE_00013539/>.

⁶⁸ fn (65), art 50.

⁶⁹ *Ibid*, art 13.

commercial speakers to make factual and uncontroversial disclosures where the requirement is reasonably related to preventing consumer deception.⁷⁰ *Zauderer* establishes a standard for narrow mandatory disclosures of objective facts, but its limits are significant. The Supreme Court's subsequent decision in *NIFLA v. Becerra* confirmed that *Zauderer* only covers compelled disclosures that are factual and non-controversial. It does not permit the government to require speakers to express broader or opinion-based messages.⁷¹ The disclosure obligation in Article 50, that users are informed they are interacting with an AI, is close to a simple factual notice and would likely survive *Zauderer* review. A U.S. statute requiring equivalent disclosure would be constitutionally defensible.

However, Article 13 goes further. Its requirement that providers characterize the accuracy, performance characteristics, and limitations of high-risk AI systems is not a simple factual disclosure.⁷² It requires providers to produce interpretive assessments and disclose those of systems whose internal workings are partially opaque⁷³ even to those who designed them, calibrated to the EU's risk-classification framework and the EU's judgment about what information deployers need to assess AI outputs. The characterisation of a system's limitations necessarily involves technical and normative judgments that are contestable rather than factual in the *Zauderer* sense. Courts applying *NIFLA* to a domestic equivalent would ask whether the mandated disclosure is ideologically freighted, technically disputed, or carries normative implications. Article 13's depth of required characterisation, and the underlying judgement of what might be harmful, suggests that stricter scrutiny would apply. Under intermediate or strict scrutiny, the government would need to demonstrate that the specific informational burden is substantially related to, or necessary for, a compelling interest. That burden may be satisfiable for some Article 13 requirements but is far from automatic for the entire clause. The constitutional defensibility of a domestic equivalent would depend on the formulation of the compelled disclosure and the government's ability to demonstrate the necessity of each element, which is a level of specificity that Article 13's general categories do not themselves provide. The First Amendment analysis of both categories reveals a common structural feature. The obligations in question are constitutionally sustainable within the EU's proportionality-based framework. They are not necessarily sustainable under the U.S. framework, which does not

⁷⁰ *Zauderer v Office of Disciplinary Counsel* 471 US 626 (1985).

⁷¹ *National Institute of Family and Life Advocates v Becerra* 585 US (2018).

⁷² fn (65), art 13(3)(b).

⁷³ Bram Vaassen, 'AI, Opacity, and Personal Autonomy' (22 September 2022) 35 *Philosophy and Technology* 88 <<https://doi.org/10.1007/s13347-022-00577-5>>, Introduction.

permit content restrictions on the ground that competing interests justify them, and which applies heightened scrutiny to compelled disclosures that extend beyond the *Zauderer* category. The implications for the analysis that follows are significant as the effects that the AI Act produces on U.S. firms' expressive conduct are constitutionally significant in ways that neither the EU legislature nor the Brussels Effect mechanism was designed to address.

V. How the AI Act shapes U.S. AI Governance in practice

1. Market access pressure and the non-divisibility dynamic

a. Why global standardisation is the rational commercial response

The preceding chapter established that the AI Act's core obligations, if enacted by Congress, would raise serious First Amendment difficulties. It is necessary to establish how significantly those obligations are already shaping U.S. AI governance in practice. It is the scale of that shaping that makes the judicial gap consequential and that creates the conditions for the legislative pressure this thesis identifies. The question is not only theoretical, but it is significant to understand how U.S. AI providers in fact respond to the market access conditions the AI Act creates, and what their behaviour reveals about the dynamics that the non-divisibility argument predicts. The answer, as far as the available evidence permits, is neither uniform adoption nor systematic resistance.

The commercial logic favouring global standardisation is most visible at the frontier of general-purpose AI development, where the firms most exposed to AI Act compliance pressure are also those for whom EU market access is commercially indispensable. The EU represents over 450 million potential users and a digital economy of considerable scale.⁷⁴ For a general-purpose AI provider whose business model depends on broad user adoption across multiple markets, withdrawal from the EU market is not a commercially rational option in the way that

⁷⁴ Eurostat, 'Population on 1 January' (2026) <<https://doi.org/10.2908/TPS00001>>.

withdrawal from a smaller or less affluent market might be. The cost of non-compliance turnover under Article 99 of the AI Act⁷⁵ combined with the reputational consequence of market exclusion, creates strong incentives to comply⁷⁶. Where compliance with the AI Act does not require conduct that violates U.S. law, the rational commercial response is to build EU compliance into the product architecture from the outset and extend it globally, rather than bear the cost of maintaining parallel systems. This dynamic is observable in the engagement of leading U.S. providers with the GPAI, which was finalised in July 2025 and addresses transparency, copyright, and systemic risk obligations for GPAI model providers under Chapter V of the AI Act. Anthropic, Google, Microsoft, and OpenAI are among its signatories.⁷⁷ Their participation is not legally required as the Code is a voluntary instrument, but it does represent a deliberate choice to engage with EU regulatory expectations at the governance and design level before the full force of the Act's high-risk obligations under Articles 9 through 15 becomes applicable in August 2026. This pattern is consistent with what Bradford identifies as the mechanism through which global standardisation occurs. Firms do not wait for enforcement to begin before internalising regulatory requirements, but do so in anticipation, because early compliance investment is cheaper than reactive restructuring.⁷⁸ The governance choices being made now about data documentation, transparency disclosures, and risk classification are shaping product architectures that will serve both EU and U.S. users through the same unified system.

However, engaging with the GPAI Code does not constitute compliance with the AI Act's most demanding obligations. The Code addresses Chapter V obligations for general-purpose models⁷⁹ especially set out in articles 53 and 55 of the AI Act, not the high-risk system obligations of Articles 9 through 15, which require structural redesign of entire AI architectures.⁸⁰ The current pattern of engagement therefore reflects selective voluntary compliance with the less architecturally demanding tier of the Act's requirements, with the deeper obligations still materialising. The significance of what is already observable is that even

⁷⁵ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, art 99.

⁷⁶ Charlotte Siegmann and Markus Anderljung, 'The Brussels Effect and Artificial Intelligence: How EU Regulation Will Impact the Global AI Market' (Governance of AI Research Paper, 16 August 2022) <https://cdn.governance.ai/Brussels_Effect_GovAI.pdf> p 29.

⁷⁷ European Commission, 'The General-Purpose AI Code of Practice' <<https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>>.

⁷⁸ Anu Bradford, 'The Brussels Effect' in *The Brussels Effect: How the European Union Rules the World* (1st edn, Oxford University Press 2020) p 42-43.

⁷⁹ European Commission, 'The General-Purpose AI Code of Practice' <<https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai#ecl-inpage-the-3-chapters-of-the-code>>.

⁸⁰ fn (75) arts 53 and 55.

at this early stage of the compliance cycle, U.S. providers are making design and governance choices calibrated to EU standards and applying them globally. The constitutional effects this produces on U.S. users are therefore not merely hypothetical but to a meaningful degree operational.⁸¹ The dynamic is not a distinctively European phenomenon but a consequence of market size and regulatory stringency operating together, as Chapter II established.

b. The GDPR comparison: why AI Act compliance is architecturally deeper

Having established that the AI Act's reach is structurally analogous to the GDPR's but architecturally deeper in consequence, that contrast deserves more precise development here, because the degree of architectural depth determines the real-world intensity of the global standardisation pressure, and it also explains why the market access-driven compliance dynamic operates differently depending on how AI products are delivered. When the GDPR came into force in 2018, U.S. firms subject to its requirements were required to adjust data handling, consent, and transfer practices.⁸² These adjustments were organisationally significant and for some cases costly, but they were largely separable from the core product architecture. A firm could modify its data collection interfaces, update its privacy notices, appoint a Data Protection Officer⁸³, and establish transfer mechanisms⁸⁴ without redesigning the underlying technology that its service ran on. The product would remain structurally the same with a compliance layer added around it. This separability is what made GDPR compliance relatively tractable for global standardisation purposes.⁸⁵ The less demanding standard applicable outside the EU could be satisfied by a product designed to meet the GDPR, because what the GDPR required was not linked to how the product worked but to how data associated with its use was handled in relation to technical and organizational measures.⁸⁶

AI Act compliance for high-risk systems under Articles 9 - 15 does not permit this separation. For example, Article 9 requires providers to establish a risk management system

⁸¹ It is worth noting that, if the U.S. were to, hypothetically, enact comparably stringent AI regulation, it would probably reach EU firms serving the U.S. market through the same market access logic.

⁸² 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) [2016] OJ L 119/1, art 3(2).

⁸³ *Ibid*, art 37.

⁸⁴ *Ibid*, art 15.

⁸⁵ Charlotte Siegmann and Markus Anderljung, 'The Brussels Effect and Artificial Intelligence: How EU Regulation Will Impact the Global AI Market' (Governance of AI Research Paper, 16 August 2022) <https://cdn.governance.ai/Brussels_Effect_GovAI.pdf> p 73.

⁸⁶ fn (82), art 25.

integrated into the entire product lifecycle, not appended to it⁸⁷ and Article 14 requires that the system be designed to enable human oversight and intervention in ways that may constrain the autonomous operation of AI systems in high-stakes domains.⁸⁸ These are not compliance layers that can be added to an existing product. They are architectural constraints that must be built into the product from the design stage. A high-risk AI system trained without satisfying Article 10's data governance requirements⁸⁹ cannot become compliant by retrospective adjustment, as the training data is or was already embedded in the model weights. Compliance requires a different training pipeline, not a different policy document. This architectural depth has a direct consequence for the market segmentation question. Under the GDPR's separable compliance model, a product built to satisfy GDPR requirements would typically find that a single product version satisfied both EU and non-EU markets, because the GDPR's requirements did not make the product less useful or less available outside the EU. Under the AI Act's embedded compliance model, a product built to satisfy Articles 9 through 15 may differ in architecturally significant ways from a product that has not been so constrained. A system designed to enable human oversight under Article 14 may be less capable of certain autonomous applications than a system designed without that constraint. A system trained to Article 10's data governance standards may differ in its statistical properties from one trained without those requirements. These differences create, at least in principle, a genuine architectural divergence between EU-compliant and non-compliant versions that has no real analogue in the GDPR context.

This distinction also explains the commercial choices of firms that have decided not to bring certain products to the EU market rather than redesign them for compliance. The decision not to release the AI features of its Ray-Ban smart glasses in the EU market is a significant illustration.⁹⁰ The product in question embeds AI capabilities into a physical device rather than delivering them through an online interface. A cloud service can be updated, retrained, and redesigned iteratively. A hardware product containing embedded AI cannot be retroactively retrained once manufactured and distributed. The combination of embedded AI, fixed hardware, and architecturally demanding compliance requirements creates a situation in which compliance cost is not fixed on an existing product but requires a fundamentally different product design from the outset. When that cost exceeds the expected revenue from the EU

⁸⁷ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, art 9(1)-(2).

⁸⁸ *Ibid*, art 14.

⁸⁹ *Ibid*, art 10.

⁹⁰ James Pero, 'Europe Can't Get Meta Ray-Ban Display Because of EU Regulations' *Gizmodo* (25 March 2026) <<https://gizmodo.com/privacy-isnt-the-only-thing-plaguing-metas-smart-glasses-2000737911>>.

market, the market exit or absence becomes the rational commercial choice, and the non-divisibility logic that drives global standardisation breaks down entirely.⁹¹

The difference between online AI services and AI built into hardware reflects a broader difference between cases where global standardisation works and cases where it does not. For general-purpose AI systems delivered over the internet, where the same model is used worldwide through one central system, this non-divisibility logic applies very strongly. Building EU compliance into the model architecture and extending it globally is cheaper than maintaining parallel systems. For hardware-embedded AI like Meta-glasses, where the delivery mechanism is physical and the product lifecycle is fixed at manufacture, the calculations seem to be done differently. The result is selective market presence rather than uniform global standardisation. Some products are redesigned for compliance and offered globally, others are withheld from the EU market entirely, and the market access-driven compliance dynamic operates differently across these product categories. This selectivity matters for the constitutional analysis, because the U.S. users most affected by globally extended AI Act compliance are precisely those using general-purpose AI systems, where the standardisation pressure is strongest and the constitutional effects most thoroughly embedded. The absence of any domestic constitutional accounting for those effects in the U.S. is therefore most acute precisely where the practical impact on U.S. expressive freedom is greatest.

2. The compliance trajectory of U.S. firms

a. Signals from early engagement: the GPAI Code of Practice

Having established that leading U.S. AI providers signed the General-Purpose AI Code of Practice, this engagement reflects anticipatory compliance rather than reactive adjustment. What that analysis left open is the more precise question of what GPAI Code engagement requires firms to do at the operational level, and what their responses reveal about the likely trajectory when the more architecturally demanding obligations of Articles 9 through 15 enter into force in August 2026. The Code was structured around four working groups addressing

⁹¹ Marco Almada and Anca Radu, 'The Brussels Side-Effect: How the AI Act Can Reduce the Global Reach of EU Policy' (2024) 25 German Law Journal 657, p. 657.

transparency of training data, copyright compliance procedures, the identification and mitigation of systemic risks, and safety and security measures for frontier models.⁹² For each of these areas, signatories commit to implementing specific governance measures within defined timeframes, to reporting on their implementation to the European AI Office, and to participating in the iterative refinement of the Code's standards as regulatory experience accumulates.⁹³ These commitments show that the involvement of the major AI companies which are signatories is not merely declaratory. In fact, they require firms to establish internal governance structures such as documentation protocols, risk assessment procedures and incident reporting channels that did not previously exist in any legally structured form under U.S. domestic law. A firm that implements these structures for the purpose of GPAI Code compliance is building organisational capacities which are directly relevant to the deeper obligations of Articles 9 through 15, even if it does not yet satisfy them. The compliance trajectory is therefore not only about whether firms will choose to comply with the AI Act's high-risk obligations when they fall due. It is about whether the organisational infrastructure being built now in response to lighter obligations creates the basis for deeper compliance or the extent to which those deeper obligations can be met.

The Anthropic case illustrates both dimensions of this dynamic. Anthropic is a GPAI Code signatory and has publicly committed to engaging with EU regulatory processes.⁹⁴ In May 2026, however, it released its Claude Mythos model, a significantly more capable system than its predecessors, shared initially with selected major U.S. technology companies and financial institutions before any public or European release.⁹⁵ The European Commission sought a meeting with Anthropic regarding the model's capabilities and its potential classification under the AI Act's high-risk framework, but Anthropic was unable to attend.⁹⁶ This episode is instructive as it illustrates the structural tension between the pace of AI capability development and the pace of regulatory engagement. The AI Act's classification framework, including the question of whether a system of Mythos's capability level could trigger high-risk obligations

⁹² European Commission, 'Second Draft of the General-Purpose AI Code of Practice Published' (19 December 2024) <<https://digital-strategy.ec.europa.eu/en/library/second-draft-general-purpose-ai-code-practice-published-written-independent-experts>>.

⁹³ European Commission, 'Signatory Taskforce for the GPAI Code of Practice' <<https://digital-strategy.ec.europa.eu/en/policies/signatory-taskforce-gpai-code-practice>>.

⁹⁴ European Commission, 'The General-Purpose AI Code of Practice' <<https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>>.

⁹⁵ Cade Metz and Kate Conger, 'Anthropic Faces Scrutiny Over Claude Mythos' *The New York Times* (12 May 2026) <<https://www.nytimes.com/2026/05/12/technology/anthropic-claude-mythos.html>>.

⁹⁶ Sam Clark, Pieter Haeck and Kathryn Carlson, 'EU Pressure Builds on Anthropic Over Mythos Hacking Risks' *Politico* (5 May 2026) <<https://www.politico.eu/article/eu-pressure-builds-on-anthropic-over-mythos-hacking-risks/>>.

under Annex III⁹⁷, was designed to operate on a regulatory timeline that the current rate of frontier AI development is already straining. The compliance trajectory is therefore shaped not only by firms' willingness to engage with EU standards but by whether the regulatory framework can keep pace with the systems it is designed to govern.

What the GPAI Code engagement demonstrates, taken together, is that the compliance trajectory of leading U.S. AI providers is real but uneven. Engagement appears to be genuine at the lighter, more procedural tier of the Act's requirements. It is anticipatory but untested at the more architecturally demanding side and it is already revealing points of friction that are likely to shape how the August 2026 deadline plays out in practice. The constitutional implications identified intensify as this trajectory advances, because each step toward deeper compliance embeds EU regulatory choices more thoroughly into the product architectures that serve U.S. users.

b. The phased timeline and the compliance pressure still materialising

Understanding the compliance steps requires situating current firm behaviour within the AI Act's phased implementation schedule, because the obligations generating the most significant architectural pressure have not yet entered into force for most providers. The timeline in article 113 of the AI Act stretches over August 2025, August 2026 and August 2027.⁹⁸ Changes have been made though with some deadlines for certain types of AI being pushed back to December 2027.⁹⁹ The current moment sits at the precise threshold of the most architecturally demanding phase of this schedule, which means that the compliance decisions most consequential for the constitutional dynamics this thesis examines are only now being made. This timing is significant for two reasons. First, it means that the compliance behaviour observable so far with the GPAI Code engagement for example, reflects adaptation to the lighter procedural tier of the Act's requirements. The harder and more revealing question of whether U.S. firms will restructure their training pipelines, documentation systems, and human oversight architectures to satisfy high-risk obligations for AI systems deployed in education,

⁹⁷ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, Annex III.

⁹⁸ *Ibid*, art 113.

⁹⁹ European Commission, 'Guidelines for providers and deployers of AI-high risk systems' <<https://digital-strategy.ec.europa.eu/en/policies/guidelines-ai-high-risk-systems>>.

employment, credit assessment, law enforcement, and the administration of justice has not yet been definitively answered by market behaviour. Some legal deadlines arrived, but enforcement decisions by national market surveillance authorities and the European AI Office will take time to develop, and firms are calibrating their responses to both the legal obligation and the realistic enforcement timeline simultaneously.¹⁰⁰ Second, the phased timeline means that the constitutional effects identified in this thesis are themselves also phased. The effects already operational from GPAI Code compliance extending EU transparency to safety standards globally are real but relatively contained. The effects that will flow from high-risk compliance decisions under Articles 9 - 15, if those decisions result in global standardisation of training data governance, human oversight architectures, and accuracy documentation requirements, will be considerably more pervasive. AI systems operating in the high-risk domains, employing screening, credit assessment, educational evaluation or law enforcement support, are precisely the systems whose outputs carry the greatest consequences for individual U.S. users. In the scenario that those systems are designed to EU high-risk standards and deployed globally through unified architectures, the expressive and procedural parameters shaping their outputs will have been set by EU regulatory choices rather than U.S. constitutional ones¹⁰¹, at a scale and in domains of far greater individual consequence than general-purpose AI.

Based on this, AI Act compliance demands may produce a structural fragmentation rather than the uniform global standardisation a non-divisibility logic would predict, with firms for whom compliance cost is manageable standardising globally and firms for whom it is prohibitive either exiting the EU market or maintaining two architecturally distinct product versions.¹⁰² Similarly the Act's high-risk obligations create differentiated compliance burdens across firm sizes and product categories, with smaller providers and more specialised AI applications facing proportionally greater restructuring costs than experienced model developers with existing compliance infrastructure.¹⁰³ Even if article 62 of the AI Act provides preferential treatment for smaller enterprises¹⁰⁴, there is no guarantee of seamless implementation by all firms willing to fully follow it and face the costs of potentially

¹⁰⁰ European Parliamentary Research Service, 'Enforcement of the AI Act' (18 March 2026)

<<https://epthinktank.eu/2026/03/18/enforcement-of-the-ai-act/>>.

¹⁰¹ Charlotte Siegmund and Markus Anderljung, 'The Brussels Effect and Artificial Intelligence: How EU Regulation Will Impact the Global AI Market' (Governance of AI Research Paper, 16 August 2022)

<https://cdn.governance.ai/Brussels_Effect_GovAI.pdf> p 38.

¹⁰² Marco Almada and Anca Radu, 'The Brussels Side-Effect: How the AI Act Can Reduce the Global Reach of EU Policy' (2024) 25 German Law Journal 657, p 657.

¹⁰³ European Parliamentary Research Service, *Implementation and Enforcement of the AI Act* (October 2025)

<[https://www.europarl.europa.eu/RegData/etudes/STUD/2025/778575/ECTI_STU\(2025\)778575_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2025/778575/ECTI_STU(2025)778575_EN.pdf)> p 24.

¹⁰⁴ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, art 62.

restructuring their AI system. The trajectory is therefore not a single curve but a distribution, as different firms, in different product categories, facing different cost-revenue calculations, will arrive at different compliance decisions over the coming months and years. The constitutional implications of that distribution are concentrated among the firms and product categories where global standardisation is the rational outcome.

3. Political resistance and its structural limits

The compliance trajectory described in the preceding two sections does not unfold in a politically neutral environment either. The current U.S. federal administration has adopted an explicitly divergent approach to AI governance, framing EU regulatory standards not as a model to be considered but as a possible obstacle to U.S. technological leadership. Understanding the structural limits of that political resistance is essential, because the market access-driven compliance dynamic continues to operate regardless of domestic political opposition, and that claim requires precise elaboration. The Trump administration's executive order of January 2025, revoking the Biden administration's AI executive order of October 2023, established the foundational posture of current U.S. AI policy, which includes the removal of regulatory barriers to AI development, the promotion of U.S. AI dominance globally, and the explicit rejection of regulatory frameworks perceived as inhibiting innovation.¹⁰⁵ This was followed in July 2025 by a further executive order directing the promotion of the export of the American AI technology stack, framing U.S. AI governance as an instrument of geopolitical competition rather than regulatory convergence.¹⁰⁶ The March 2026 National Policy Framework for AI elaborated this approach at the legislative level, recommending that Congress adopt a single national standard for AI governance designed to pre-empt state-level regulation and to resist the adoption of more stringent foreign frameworks.¹⁰⁷ Taken together, these developments represent a coherent and sustained political effort to ensure that U.S. AI governance develops

¹⁰⁵ Executive Office of the President, 'Removing Barriers to American Leadership in Artificial Intelligence' (White House, 23 January 2025) <<https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/>>.

¹⁰⁶ Executive Office of the President, 'Promoting the Export of the American AI Technology Stack' (White House, 23 July 2025) <<https://www.whitehouse.gov/presidential-actions/2025/07/promoting-the-export-of-the-american-ai-technology-stack/>> section 1.

¹⁰⁷ Executive Office of the President, *National Policy Framework for Artificial Intelligence: Legislative Recommendations* (March 2026) <<https://www.whitehouse.gov/wp-content/uploads/2026/03/03.20.26-National-Policy-Framework-for-Artificial-Intelligence-Legislative-Recommendations.pdf>>.

independently of EU regulatory standards. The central question is whether that political effort can dissolve the compliance pressure that the AI Act's market access conditionality creates for U.S. firms. The answer is that it cannot, for reasons that are structural rather than political.

The executive orders and legislative recommendations operate exclusively within U.S. domestic law. They shape what the U.S. government requires of AI providers operating domestically and cannot alter what EU law requires of providers that wish to operate in the EU market. The obligation imposed by Article 2 of the AI Act on providers placing AI systems on the Union market, and the enforcement regime established by Articles 22 and 99, derive from EU law and are entirely independent of U.S. domestic policy preferences.¹⁰⁸ A U.S. firm that wishes to serve EU users must comply with the AI Act regardless of what the U.S. executive branch recommends, because the legal obligation is not one that U.S. domestic authority can waive. Political divergence between the U.S. and EU approaches to AI governance therefore does not eliminate the compliance incentive that drives global standardisation. It complicates the institutional environment in which compliance decisions are made, without altering the fundamental market access calculus that makes compliance the rational commercial choice for firms with significant EU exposure.

The more significant consequence of political divergence lies elsewhere. A federal government that actively frames EU regulatory standards as incompatible with U.S. values and that signals institutional resistance to any regulatory convergence reduces the political acceptability of treating EU compliance as a model for domestic operations. Firms operating in both markets may be less inclined to present their EU compliance investments as governance best practices applicable domestically, and more inclined to characterise them as market-specific legal obligations that do not reflect their preferred operational standards. This characterisation matters for the constitutional analysis. It reinforces the voluntary character of the global standardisation decision, because it makes explicit that firms are extending EU standards globally as a commercial choice rather than as a reflection of any normative alignment with EU regulatory values. The state action requirement identified as the central obstacle to constitutional review remains absent, and the political context turns out to sharpen rather than soften the constitutional paradox. Firms are applying EU speech standards to U.S. users not

¹⁰⁸ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, arts 2, 22 and 99.

because they believe those standards are constitutionally appropriate domestically, but because the economics of market access make it cheaper to do so than to maintain separate architectures.

There is a further dimension to the fragmentation scenario that political resistance introduces. If sustained U.S. regulatory divergence from EU standards eventually makes the compliance cost of maintaining a single globally standardised product architecture high, because EU and U.S. requirements diverge so sharply that satisfying both simultaneously becomes architecturally impossible, the non-divisibility logic breaks down and market segmentation becomes unavoidable.¹⁰⁹ This scenario, in which firms are compelled to maintain EU-compliant and U.S.-specific product versions, would represent the structural limit of the market access-driven compliance dynamic in the AI sector. It would also represent the point at which the constitutional tension dissipates not through resolution but through architectural separation. In that sense, EU regulatory choices would no longer shape the AI systems serving U.S. users, because those systems would no longer be the same systems serving EU users. The constitutional problem disappears, but so does the economic logic for the global standardisation that made European digital market regulation globally consequential in the first place. Whether political divergence reaches that threshold depends on the degree of regulatory incompatibility that develops in the meantime of the enforcement of the entire AI Act, the conditions required for the market access-driven compliance dynamic to operate in the field of AI, the fact that EU standards are stricter than U.S. ones and simultaneous compliance being possible. Political resistance and the structural limits are clearly present. It is within those limits that the constitutional tension continues to accumulate, unresolved by courts and yet unaddressed by Congress.

¹⁰⁹ Marco Almada and Anca Radu, 'The Brussels Side-Effect: How the AI Act Can Reduce the Global Reach of EU Policy' (2024) 25 German Law Journal 657, p 657.

VI. The Failure of Judicial Remedies: Foreign Regulatory Diffusion and the Limits to Constitutional Litigation

1. Why the Judicial Gap Is Structural and Not Incidental

a. Voluntary Compliance and the Absence of Constitutional Review

Overall, the First Amendment constrains government actors, and it does not regulate the expressive choices of private firms. When a U.S. AI provider extends EU compliance obligations to its domestic operations, such as restricting categories of output that Article 5 prohibits, or generating disclosures that Articles 13 and 50 require, it does so as a private commercial actor making a voluntary business decision. As established, no provision of U.S. law compels it to do so and no government actor has ordered, incentivised, or mandated that choice. The constitutional protections whose potential violation the previous sections identified are therefore, in the ordinary case, structurally unavailable. The First Amendment's reach is bounded by the state action requirement, which limits constitutional obligations to governmental actors and their agents. The requirement is not a procedural but a foundational principle of U.S. constitutional architecture. The Constitution's guarantees of individual rights are, with limited exceptions, guarantees against government, not against the world.¹¹⁰ A private party's decision to restrict speech, however consequential for the persons affected, does not become a constitutional violation simply because the restriction resembles one that the government could not impose. The question is not whether effects on speech produced by private AI providers are significant. It is whether a government actor is responsible for that effect. In the context of compliance with the AI Act for private companies, the issue of state action is usually clear and does not support a constitutional claim. A U.S. AI provider that implements Article 5 prohibitions globally, declining to generate categories of output for its domestic users that the EU legislature has determined are incompatible with its definition of fundamental rights, is exercising its own editorial discretion.¹¹¹ As the Supreme Court confirmed in *Moody v. NetChoice*, private platforms' decisions about what content their systems carry, rank, or suppress are themselves protected expressive activity under the First

¹¹⁰ Frank I Michelman, 'The State Action Doctrine' in Vikram D Amar and Mark V Tushnet (eds), *Global Perspectives on Constitutional Law* (Oxford University Press 2009) 228, p 228-229.

¹¹¹ Evelyn Douek and Genevieve Lakier, 'Rereading "Editorial Discretion"' (Knight First Amendment Institute, 24 October 2022) <<https://knightcolumbia.org/blog/rereading-editorial-discretion>>.

Amendment.¹¹² The provider is not being compelled to restrict speech but it is choosing to do so, and that choice is constitutionally protected rather than constitutionally problematic from the provider's own perspective. The users whose access to certain AI-generated content is thereby curtailed have no constitutional claim against the provider, because the provider is a private actor. They have no constitutional claim against the U.S. government, because the U.S. government has done nothing. And they have no actionable claim against the EU, because EU law does not govern U.S. constitutional rights.

This point requires precise statement because it is frequently concealed by the language of regulatory pressure or market compulsion. The non-divisibility dynamic established in Chapter II describes an economic constraint, with firms finding it rational to maintain a single globally compliant product rather than parallel market-specific versions. But economic rationality is not legal compulsion. The Constitution does not mandate that private actors optimise for constitutional values in their commercial decisions, it mandates only that government actors respect constitutional limits when they exercise their authority.¹¹³ A firm that extends EU standards globally because doing so is cheaper than segmenting its operations is making a business judgment, not responding to a government mandate. The constitutional framework has nothing to say about that judgment directly, however significant its effects on expression may be for the users or the companies. Putting this dynamic in the context of the Brussels Effect, the de facto Brussels Effect operates because firms act voluntarily. EU compliance is extended not because any external authority outside the EU requires firms to do so, but because the economics of market access make global standardisation the least-cost strategy.¹¹⁴ This voluntary character is what distinguishes the de facto from the de jure mechanism, in which foreign governments actively adopt EU standards into their own legal frameworks. In the de facto case, no sovereign legislative act occurs in the receiving jurisdiction. No government actor in the U.S. decides that AI Act standards should govern domestic conduct. The diffusion occurs entirely through private corporate decisions that are, from a U.S. constitutional standpoint, the exercise of the providers' own rights rather than an infringement of anyone else's.

The doctrinal consequence is significant. There is no defendant against whom a First Amendment claim arising from globally applied AI Act compliance obligations can be brought

¹¹² *Moody v NetChoice LLC*, 603 US(2024) No 22–277, Opinion of the Court p 14.

¹¹³ *Blum v Yaretsky*, 457 US 991 (1982).

¹¹⁴ Anu Bradford, 'The Brussels Effect' in *The Brussels Effect: How the European Union Rules the World* (1st edn, Oxford University Press 2020) p 28.

in the Supreme Court. The U.S. government has not acted. Even if the AI provider has acted, its action is constitutionally protected. The EU has acted, but its action is not reviewable by U.S. courts applying U.S. constitutional standards. Analysing the implications of *Moody v. NetChoice* for platform regulation, the First Amendment protects private editorial decisions from government interference, but it does not provide a way to limit the broader effects those decisions may have on public debate, even when they are influenced by foreign regulation.¹¹⁵ The constitutional framework is built around the idea that the government and private parties are distinct actors. When the most consequential shaping of speech occurs through private commercial decisions responding to foreign regulatory incentives, that framework cannot reach it.

One potential exception deserves consideration. The state action doctrine recognises that nominally private conduct may be attributed to the government where there is sufficient entanglement between the private actor and a government authority. The Supreme Court has identified several theories of state action, including the public function doctrine, the joint participation doctrine, and the nexus or symbiosis test, each of which locates the conditions under which private conduct may be treated as governmental for constitutional purposes.¹¹⁶ In the AI Act compliance context, state entanglement might only arise where a U.S. government agency deploys an EU-compliant AI system in a context affecting legally protected interests of U.S. persons or where a regulatory framework creates sufficiently close institutional ties between government authority and private compliance conduct. Were such entanglement demonstrated, the First Amendment analysis developed before would become directly operative. These scenarios of governmental entanglement, while legally possible, are peripheral to the mechanism this thesis examines. The Brussels Effect operates through private commercial decisions made at scale. The constitutional tension it generates is therefore systemic and not case-specific to any government act. A doctrine designed to identify government responsibility in specific transactions is structurally not suited to address effects that arise from the compiled operation of private compliance decisions across an entire market. The state action doctrine reflects a historical moment in which the primary threat to constitutional rights was direct government coercion. It was not designed for regulatory environments in which the most consequential shaping of protected conduct occurs through

¹¹⁵ Jack M Balkin, 'Moody v NetChoice: The Supreme Court Meets the Free Speech Triangle' (Yale Law School Public Law Research Paper, 19 January 2025) <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5104013> p 34.

¹¹⁶ *Burton v Wilmington Parking Authority*, 365 US 715 (1961).

private actors responding to incentives, whether domestic or foreign.¹¹⁷ The result is that the constitutional tension identified in section 1 is real but judicially not addressable in most of the cases where it matters most. The next section examines what that structural unavailability means for the constitutional values at stake.

b. The Gap Between Constitutional Significance and Judicial Remedy

What remains to be established is the full significance of the structural gap of why no available plaintiff can bring the claim that the constitutional analysis seems to demand, and what happens to the constitutional tension that litigation cannot resolve. The threshold requirement for access to federal court is standing under Article III of the U.S. Constitution¹¹⁸, which the Supreme Court in *Lujan v. Defenders of Wildlife* distilled into three elements: the plaintiff must have suffered an injury in fact that is concrete and particularised, the injury must be fairly traceable to the challenged conduct, and the injury must be redressable by a favourable judicial decision.¹¹⁹ Each of these elements presents a distinct difficulty in the AI Act compliance context.

The injury requirement demands that the plaintiff identify a concrete harm that has materialised, not a speculative or abstract concern about regulatory diffusion. A U.S. user whose access to certain categories of AI-generated content is restricted because a provider has extended Article 5 prohibitions globally might argue that they have suffered a concrete expressive injury. But that injury, if it exists, flows from the provider's editorial decision, not from any government act. The provider is exercising its own discretion, and the user has no constitutional entitlement to compel the provider to generate content it has chosen not to generate.¹²⁰ The First Amendment protects private editorial decisions, even when those decisions exclude others from the platform or space being controlled. This exclusion is considered a constitutional feature, not a constitutional problem.¹²¹ The user's injury, such as it

¹¹⁷ Jed Rubenfeld, 'The First Amendment's Purpose' (April 2001) 53 *Stanford Law Review* 767 <https://heinonline-org.uaccess.univie.ac.at/HOL/Page?collection=journals&handle=hein.journals/stflr53&id=786&men_tab=srchresults> p 793-94.

¹¹⁸ US Constitution art III.

¹¹⁹ *Lujan v Defenders of Wildlife*, 504 US 555 (1992).

¹²⁰ *Simon v Eastern Kentucky Welfare Rights Organization*, 426 US 26 (1976).

¹²¹ Frederick Schauer, 'Towards an Institutional First Amendment' (18 February 2005) 89 *Minnesota Law Review* 1256 <<https://scholarship.law.umn.edu/cgi/viewcontent.cgi?article=1681&context=mlr>> p 1275.

is, is not a constitutional one because no government actor has caused it, even if he claims that constitutional rights are involved.

The traceability requirement adds to the difficulty. Even if a plaintiff could identify a cognizable injury, they would need to demonstrate that it is traceable to the conduct of a government actor rather than to the independent decision of a private firm. Where a U.S. AI provider voluntarily extends EU compliance standards to its domestic operations, the causal chain between any government act and the plaintiff's injury runs through the firm's autonomous commercial judgment. The Supreme Court's decision in *Allen v. Wright* established that traceability fails where the causal connection between government conduct and plaintiff injury depends on the independent choices of a third party not before the court.¹²² The voluntary compliance dynamic is precisely such a case. The U.S. government has not acted, and the EU's market access conditions are not directly enforceable against U.S. users. The firm's decision to globalise its compliance architecture is the operative cause of whatever expressive effects are felt domestically. No plaintiff challenging those effects in a U.S. federal court can trace them to a government actor in the manner Article III requires.

The redressability requirement adds a further layer. Even if standing were somehow established, a federal court ruling in the plaintiff's favour could not remedy the systemic dynamic that produces the constitutional concern. A court could, in theory, stop a government actor from using a specific EU compliant AI system in a particular situation. It could not order U.S. AI providers to stop extending EU compliance globally, because those providers are exercising their own constitutionally protected discretion. It could not stop the EU from setting the market access conditions that create the compliance incentive, because EU regulatory authority is not subject to U.S. judicial review. And it also could not reverse the commercial logic of non-divisibility that makes global standardisation the rational response to those conditions. Looking at the limits of judicial review in structural regulatory contexts, courts are generally institutionally equipped to resolve specific disputes between identified parties over concrete injuries but are poorly positioned to address systemic effects that arise from the action of private market decisions responding to structural incentives. The constitutional concern generated by EU regulatory diffusion is precisely of the systemic kind that courts cannot reach. The result is an unusual constitutional landscape. The obligations the AI Act imposes remain constitutionally significant in the sense that equivalent domestic legislation would face serious

¹²² *Allen v Wright* 468 US 737 (1984).

First Amendment scrutiny. The effects those obligations produce when extended globally by U.S. firms are felt by U.S. users interacting with AI systems whose expressive parameters have been calibrated to EU constitutional choices rather than U.S. ones. Those effects are real and they engage the core values that the First Amendment exists to protect. Yet no plaintiff can bring the claim that would make those effects judicially cognizable, because the state action requirement, the standing doctrine, and the limits of the remedies combine the systemic dynamic beyond the reach of constitutional litigation. The constitutional tension is present with judicial remedy being structurally absent.

This gap between constitutional significance and judicial remedy is not an accident of procedural doctrine. It reflects a certain point about the relationship between the import of foreign norms by powerful private actors and the institutional design of U.S. constitutional law. The First Amendment was constructed around a model of government as the primary threat to expressive freedom, and judicial review was designed to constrain government actors who exercise coercive authority over private speech. The Brussels Effect operates through a different model entirely. Market access conditions are set by a foreign regulator, internalised by private firms voluntarily and producing expressive effects on domestic users without any domestic government actor having exercised coercive authority at any point in the chain. The constitutional framework encounters this model at its structural limits. It can identify the values at stake but cannot provide the remedy those values seem to require. What it cannot resolve through litigation does not disappear. As the following parts establish, it accumulates as pressure in the domain where the U.S. constitutional order does retain the institutional capacity to act through the legislative domain and ultimately Congress.

2. The Conditions for Congressional Intervention

a. The First Amendment Problem as a Leftover for Legislation

The structural failure of judicial remedies does not make the First Amendment problem disappear. What courts cannot assess because state action is absent, what executive orders cannot resolve because they operate only within domestic law, and what market dynamics cannot self-correct because the economics of non-divisibility continue to favour global

standardisation, is an accumulation of a constitutional question without being able to be properly categorized institutionally. Congress is the only institution in the U.S. constitutional order with the authority to name that problem formally, assess it against First Amendment standards, and determine what domestic law should say about it. Therefore, this is where the U.S. constitutional order places the authority to make precisely this kind of determination.¹²³ The core of the problem is a democratic legitimacy deficit. When U.S. AI providers extend AI Act compliance globally through unified architectures, the content parameters and disclosure requirements governing AI systems used by U.S. users are being set by a foreign legislature applying a foreign framework. The EU legislature applied proportionality analysis calibrated to the EU Charter of Fundamental Rights¹²⁴ when it enacted Article 5's content prohibitions and Articles 13 and 50's disclosure obligations.¹²⁵ That analysis was not designed to satisfy, and was not tested against, the strict scrutiny that the First Amendment imposes on content-based restrictions¹²⁶ or the *Zauderer* standard that governs compelled commercial disclosure.¹²⁷ The result is that the expressive environment of U.S. AI users is being shaped by constitutional choices that were never subjected to U.S. constitutional review, made by a legislature with no democratic accountability to U.S. citizens, and applied through private compliance decisions that no U.S. institution has authorised or assessed. The First Amendment reflects a foundational commitment to the proposition that the boundaries of permissible expression in the U.S. are set by U.S. democratic deliberation, tested against U.S. constitutional standards, and reviewable by U.S. courts. When those boundaries are instead set by a foreign legislature operating under a different constitutional order and then adopted by private market decisions, the democratic accountability structure that the First Amendment presupposes is bypassed. Courts cannot restore it because state action is absent. Only Congress can, by enacting domestic legislation that subjects equivalent obligations to the constitutional scrutiny they have so far avoided.

What congressional action would need to accomplish is to bring the categories of AI Act obligation that produce First Amendment friction into the domestic legislative process where they can be assessed on U.S. constitutional terms. This does not require Congress to replicate or reject the AI Act as a whole. It requires Congress to make its own determination about which content restrictions on AI outputs are compatible with the First Amendment under

¹²³ Amdt1.7.13.6 Legislative Investigations (*US Constitution Annotated*, LII/Legal Information Institute) <<https://www.law.cornell.edu/constitution-conan/amendment-1/legislative-investigations>>.

¹²⁴ Charter of Fundamental Rights of the European Union [2012] OJ C 326/391, art 52.

¹²⁵ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, arts 5, 13 and 50.

¹²⁶ *Reed v Town of Gilbert*, 576 US 155 (2015).

¹²⁷ *Zauderer v. Office of Disc. Counsel*, 471 U.S. 626 (1985).

strict scrutiny, and which disclosure obligations fall within the *Zauderer* category of factual and non-controversial commercial disclosures rather than compelled interpretive speech.¹²⁸ Those are determinations that only a domestic legislature can make with constitutional authority in the U.S. legal order. Making them is what would transform the current situation, in which EU constitutional choices govern U.S. expressive conduct through private compliance decisions, into a scenario in which U.S. constitutional choices govern U.S. expressive conduct through democratically enacted and judicially reviewable domestic law. The point was captured directly during the 2023 Senate Judiciary hearings, when Senator Tillis observed that courts should not be left to provide guidance by applying existing laws to new AI scenarios, and that Congress should instead make the law clear.¹²⁹ The constitutional dimension identified gives that observation structural force. In fact, legislative clarity is not only preferable, but in the specific context of foreign regulatory diffusion through private compliance, legislative action is the only constitutionally adequate response.

b. Triggering Conditions and Available Regulatory Models

The argument that Congress is structurally positioned to act does not by itself establish when it will act or what form that action is likely to take. The triggering conditions for congressional intervention are already visible currently and they are sharpened by the constitutional dimension. Congressional attention to AI governance has been sustained since 2023, though its character has shifted significantly over that period. The May 2023 Senate Judiciary Subcommittee hearing on the oversight of AI¹³⁰, at which the CEOs of OpenAI¹³¹ and IBM¹³² testified alongside academic experts, reflected a genuine interest in federal regulation, with Chair Blumenthal and Ranking Member Hawley both emphasising the urgent need for

¹²⁸ *National Institute of Family and Life Advocates v Becerra* 585 US (2018).

¹²⁹ Crowell & Moring LLP, ‘The US Senate Hearing on Artificial Intelligence and Intellectual Property’ (6 December 2023) <<https://www.crowell.com/en/insights/client-alerts/the-us-senate-hearing-on-artificial-intelligence-and-intellectual-property>>.

¹³⁰ United States Senate Committee on the Judiciary, ‘Oversight of AI: Rules for Artificial Intelligence’ (Hearing, 16 May 2023) <<https://www.judiciary.senate.gov/committee-activity/hearings/oversight-of-ai-rules-for-artificial-intelligence>>.

¹³¹ Sam Altman, ‘Written Testimony before the Senate Judiciary Subcommittee on Privacy, Technology and the Law’ (16 May 2023) <<https://www.judiciary.senate.gov/imo/media/doc/2023-05-16%20-%20Bio%20&%20Testimony%20-%20Altman.pdf>> p 12.

¹³² Christina Montgomery, ‘Written Testimony before the Senate Judiciary Subcommittee on Privacy, Technology and the Law’ (16 May 2023) <<https://www.judiciary.senate.gov/imo/media/doc/2023-05-16%20-%20Testimony%20-%20Montgomery.pdf>> p 5.

legislative action on AI risk, transparency, and accountability. That hearing, the first in a series addressing whether and to what extent Congress should regulate AI, expressed broad consensus that federal action was necessary. The July 2023 follow-up hearing on principles for AI regulation, at which Anthropic CEO Dario Amodei testified¹³³ alongside other industry leaders, reinforced that need and focused on the substantive content of a possible federal framework.¹³⁴ The institutional climate in 2023 was such that comprehensive federal AI legislation seemed not only possible, but imminent.

By 2025 and 2026 that momentum had fractured. The May 2025 Senate Commerce Committee hearing, titled "Winning the AI Race: Strengthening U.S. Capabilities in Computing and Innovation" chaired by Senator Cruz, reframed the legislative question entirely.¹³⁵ The hearing examined how removing regulatory barriers on the AI supply chain could accelerate innovation and secure U.S. domination. The new way to compete in AI was to outpace rivals through innovation rather than to saddle AI developers with European-style regulations.¹³⁶ The shift from the 2023 framing, where regulation was seen as a necessary response to the risks AI usage might present, to the 2025 framing considering deregulation as a necessary condition of competitiveness, is institutionally significant. It means that the legislative conditions for comprehensive federal AI legislation have not simply stalled but have actively moved in a direction that makes EU-aligned regulatory models politically unavailable as a template for national action

The failure of the federal pre-emption attempts reinforces this picture of institutional fragmentation. A provision in the One Big Beautiful Bill Act that would have imposed a ten-year moratorium on state AI regulation was stripped by a near-unanimous Senate vote in July 2025 following opposition from senators, members of the House, and state governors on both sides of the aisle. A comparable provision was subsequently rejected in the context of the National Defense Authorization Act for fiscal year 2026.¹³⁷ The 99-1 vote is particularly

¹³³ Dario Amodei, 'Written Testimony before the Senate Judiciary Committee' (26 July 2023)

<https://www.judiciary.senate.gov/imo/media/doc/2023-07-26_-_testimony_-_amodei.pdf>.

¹³⁴ Jason Matheny Russell, 'Written Testimony before the Senate Judiciary Committee' (26 July 2023)

<https://www.judiciary.senate.gov/imo/media/doc/2023-07-26_-_testimony_-_russell.pdf> p 1.

¹³⁵ US Senate Committee on Commerce, Science, and Transportation, 'Winning the AI Race: Strengthening U.S. Capabilities in Computing and Innovation' (Press Release, 29 April 2025)

<<https://www.commerce.senate.gov/press/rep/release/winning-the-ai-race-strengthening-u-s-capabilities-in-computing-and-innovation-2025-4/>>.

¹³⁶ US Congress, *Winning the AI Race: Strengthening U.S. Capabilities in Computing and Innovation* (Senate Hearing, 119th Congress, 8 May 2025) <<https://www.congress.gov/119/chrg/CHRG-119shrg61426/CHRG-119shrg61426.pdf>> p 40.

¹³⁷ Ropes & Gray LLP, 'Examining the Landscape and Limitations of the Federal Push to Override State AI Regulation' (11 March 2026), section I, 'The Existing Federal and State Regulatory Frameworks for AI'

instructive because it reflects not a rejection of federal AI governance as such, but a rejection of the specific mechanism of top-down pre-emption without substantive regulatory content. Congress chose not to block state-level experimentation without creating a federal alternative, showing an awareness that the lack of AI governance might not persist¹³⁸, even though there is still disagreement about what federal regulation should look like.

The constitutional dimension identified adds a distinct and yet underarticulated dimension to that awareness. The political debate over federal AI governance has proceeded largely on the axes of innovation policy, competitive strategy, and the proper division of regulatory authority between federal and state levels. The First Amendment question has not yet entered that debate with structural force, because the mechanism through which it arises, being voluntary private compliance with foreign regulatory standards diffused, does not produce the kind of acute case or controversy that places it directly on a congressional agenda today. What will place it there is the advancing compliance trajectory. The AI Act's high-risk obligations under Articles 9 through 15 once fully applicable with the domains they govern are precisely those in which AI outputs carry the greatest individual consequences for U.S. users. As groups in those sectors begin to experience and express the concrete effects of AI systems shaped by EU rather than U.S. constitutional standards, the political importance of the First Amendment gap will grow. The Anthropic Claude Mythos episode, in which a significantly more capable AI system was released to major U.S. commercial and financial institutions before any engagement with European regulators, illustrates precisely the pace at which capability development is already generating exactly those consequences.¹³⁹

On the question of available regulatory models, content restrictions equivalent to Article 5's prohibitions would need to survive strict scrutiny, especially where Congress would need to demonstrate that each prohibited category of AI output targets a compelling governmental interest and that the restriction is the least restrictive means of serving it. Some of Article 5's categories may be capable of surviving that scrutiny if formulated with sufficient precision to avoid vagueness challenges, particularly those targeting subliminal manipulation techniques¹⁴⁰,

<https://www.ropesgray.com/en/insights/alerts/2026/03/examining-the-landscape-and-limitations-of-the-federal-push-to-override-state-ai-regulation>.

¹³⁸ McDermott Will & Emery, 'No State AI Law Moratorium in One Big Beautiful Bill Act' (5 July 2025) <https://www.mcdermottlaw.com/insights/no-state-ai-law-moratorium-in-one-big-beautiful-bill-act/>.

¹³⁹ Sam Clark, Pieter Haeck and Kathryn Carlson, 'EU Pressure Builds on Anthropic Over Mythos Hacking Risks' *Politico* (5 May 2026) <https://www.politico.eu/article/eu-pressure-builds-on-anthropic-over-mythos-hacking-risks/>.

¹⁴⁰ Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1, art 5(1).

while others dependent on broader concepts requiring case-by-case judgment would require significantly more precise legislative formulation than the AI Act itself provides. Disclosure obligations equivalent to Articles 13 or 50¹⁴¹ present a more manageable constitutional path. The basic obligation to inform users that they are interacting with an AI system falls within *Zauderer's* category of factual and non-controversial commercial disclosures.¹⁴²

The deeper obligation in Article 13 to characterise the accuracy, performance characteristics, and limitations of high-risk systems would require a formulation limited to objectively verifiable performance references, rather than normative assessments calibrated to a foreign risk classification framework, in order to remain within *Zauderer* and avoid triggering intermediate or strict scrutiny as compelled interpretive speech. As confirmed in *NIFLA v. Becerra*, the *Zauderer* doctrine does not extend to disclosures that are ideologically freighted or carry normative implications beyond the communication of objective fact¹⁴³, and Article 13's full depth of required characterisation risks crossing that line if transposed without modification. These are not insuperable obstacles. They are the constitutional constraints within which Congress would need to work, which is precisely what the EU legislature, operating under a different constitutional order, was not required to do. The point is not that the AI Act's substantive judgments are wrong but that in the case of the U.S. those judgments need to be made again, domestically, by an institution with democratic accountability to U.S. citizens and subject to review by U.S. courts applying U.S. constitutional standards. That is the institutional logic that makes congressional action not merely available but, as the compliance trajectory advances and the constitutional gap widens, increasingly necessary especially after U.S. firms show an interest in following the first obligation of the enforced AI Act of the EU.

¹⁴¹ *Ibid*, arts 13 and 50.

¹⁴² *Zauderer v Office of Disciplinary Counsel* 471 US 626 (1985).

¹⁴³ *National Institute of Family and Life Advocates v Becerra* 585 US (2018).

VII. The Constitutional Gap and the Necessity of Congressional Action

The preceding chapters have established three major findings. First, the EU AI Act's core obligations, its content prohibitions under Article 5 and disclosure requirements under Articles 13 and 50, would raise serious First Amendment difficulties if enacted by Congress in the same way, both as content-based restrictions subject to strict scrutiny and as compelled disclosures exceeding what the *Zauderer* standard permits. Second, those obligations are already shaping the product architectures of AI systems serving U.S. users through the commercial rationality of unified global compliance, with the most architecturally demanding phase of that shaping only now materialising as the implementation deadlines of the AI Act approach. Third, no judicial remedy is available to address the constitutional effects that result, because the voluntary character of private compliance decisions means that state action is absent, standing cannot be established, and redressability cannot reach the systemic dynamic that produces the constitutional concern. These three findings together generate a single institutional conclusion: the constitutional tension that courts cannot resolve and that executive action cannot dissolve must be addressed by Congress. The absence of judicial remedy does not by itself establish that Congress must act. Courts cannot review the constitutional effects of AI Act compliance because the state action requirement is unsatisfied. That is the First Amendment's architecture working as designed, not a deficiency to be overcome. The case for congressional action rests on a different ground. The First Amendment presupposes a particular accountability structure, which is that the effective limits on expression in the U.S are set by institutions democratically accountable to U.S. citizens, measured against U.S. constitutional standards, and reviewable by U.S. courts. When EU compliance obligations set those limits through private commercial decisions, the accountability structure is not violated but absent. No U.S. institution has assessed the relevant parameters against First Amendment doctrine, no domestic deliberation has authorised them and no federal court stands available to the U.S. user who believes they are wrong. Congressional action would not remedy a constitutional violation that has occurred. It could restore an accountability structure that the First Amendment presupposes and that foreign regulatory diffusion through private compliance has displaced. That is the distinction between a congressional response that is merely preferable and one that is structurally necessary. This chapter examines why that conclusion is not merely logical but increasingly politically unavoidable, and what the conditions for congressional action look like.

The most important point to understand about the constitutional gap this thesis identifies is that it is not static. It deepens as the compliance trajectory advances. The obligations already embedded in AI systems through engagement with the General-Purpose AI Code of Practice represent the lighter procedural part of what the AI Act requires.¹⁴⁴ The obligations that become fully applicable under Articles 9 through 15 are architecturally far more demanding as they govern AI systems deployed in employment screening, credit assessment, educational evaluation, law enforcement support, and the administration of justice. These are not peripheral domains, but domains in which AI outputs carry important individual consequences for users, and in which the absence of domestic constitutional accounting for how those outputs are governed is most politically and legally significant. As AI systems designed to EU high-risk standards are likely to be deployed globally through unified architectures across these domains, the expressive and procedural parameters shaping their outputs will have been calibrated to EU constitutional choices rather than U.S. ones, at a scale and in contexts of far greater individual consequence than current general-purpose AI use.

The congressional record from 2023 to 2026 reveals an institution that has registered the need for federal AI governance without yet producing the legislative response that need demands. The 2023 Senate Judiciary hearings expressed genuine urgency for legislation, and the 2025 Senate Commerce Committee reframed the question in competitive rather than regulatory terms, with the explicit rejection of an adoption of a framework similar to the AI Act, with Senators outlining the “EU’s nasty regulatory approach”.¹⁴⁵ The vote stripping the AI pre-emption provision from the One Big Beautiful Bill Act, and the comparable rejection in the National Defense Authorization Act context, demonstrate that Congress has twice declined to foreclose state-level AI regulation without offering a federal alternative yet.¹⁴⁶ That pattern shows that lawmakers know the lack of regulation cannot continue, but still cannot agree on how to address it, while other jurisdictions have already completed that step. The constitutional dimension this thesis identifies has not yet entered that debate with structural force. When and if it does, it will add a qualitatively different kind of pressure to the existing legislative dynamic but not the pressure of innovation policy or competitive strategy, but one of democratic

¹⁴⁴ European Commission, ‘The General-Purpose AI Code of Practice’ <<https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>>.

¹⁴⁵ US Congress, *Winning the AI Race: Strengthening U.S. Capabilities in Computing and Innovation* (Senate Hearing, 119th Congress, 8 May 2025) <<https://www.congress.gov/119/chrg/CHRG-119shrg61426/CHRG-119shrg61426.pdf>> p 2.

¹⁴⁶ Ropes & Gray LLP, ‘Examining the Landscape and Limitations of the Federal Push to Override State AI Regulation’ (11 March 2026), section V, ‘Looking Ahead’ <<https://www.ropesgray.com/en/insights/alerts/2026/03/examining-the-landscape-and-limitations-of-the-federal-push-to-override-state-ai-regulation>>.

accountability. The argument will not be that Congress should regulate AI because the EU has done so, but that the First Amendment requires domestic deliberation over the boundaries of permissible AI-mediated expression, and that the absence of such deliberation is itself a constitutional deficit that deepens as firms' international compliance with frameworks such as the AI Act advances.

What congressional action would need to accomplish is precise. It does not require comprehensive AI regulation modelled on the AI Act or any other foreign framework. In fact, it requires Congress to bring the specific categories of obligation that produce First Amendment friction, specifically content restrictions on AI outputs and compelled disclosure requirements, into the domestic legislative process where they can be assessed against U.S. constitutional standards. Content restrictions equivalent to the AI Act's Article 5 prohibitions would need to survive strict scrutiny, with Congress demonstrating that each prohibited category targets a compelling governmental interest, an outcome that is unlikely to hold for all of Article 5's prohibitions. Some categories might be capable of satisfying that standard. Others, dependent on broader concepts requiring case-by-case judgment, would require significantly more precise legislative formulation to avoid vagueness challenges, which are disqualifying for regulation touching upon the freedom of speech. Disclosure obligations present a more tractable path. The basic requirement that users be informed they are interacting with an AI system falls within *Zauderer's* category of factual and non-controversial commercial disclosure.¹⁴⁷ Deeper requirements characterising system limitations and performance would need to be formulated around objectively verifiable metrics rather than normative assessments calibrated to EU risk classification frameworks, in order to remain within *Zauderer* and avoid the scrutiny that *NIFLA v. Becerra* reserves for ideologically freighted or normatively loaded compelled speech.¹⁴⁸ None of these are impossible or rigid constitutional obstacles. They are the constraints within which domestic AI legislation must work, and working within them is precisely what transforms the current situation, where EU constitutional choices govern U.S. expressive conduct through private compliance decisions into one where U.S. constitutional choices do so through democratically enacted and judicially reviewable law.

The case for congressional action is therefore not that the EU AI Act is wrong, or that its regulatory judgments are misguided. Those judgments were made within a constitutional

¹⁴⁷ *Zauderer v Office of Disciplinary Counsel* 471 US 626 (1985).

¹⁴⁸ *National Institute of Family and Life Advocates v Becerra* 585 US (2018).

order whose balancing framework is well-suited to making them. The case is that the same judgments have not been made within the U.S. constitutional order, and that the mechanism through which they are currently reaching U.S. users bypasses classical accountability structures the First Amendment presupposes. It is here that the concept of the Brussels Effect acquires its full analytical significance for this thesis, proving the capacity of EU regulatory standards to achieve global reach through the private compliance decisions of firms that cannot commercially afford to forgo EU market access in an AI context. That capacity is real, and the AI Act is already exercising it now. But what the Brussels Effect framework does not fully anticipate is the constitutional resistance it encounters in jurisdictions whose legal orders are built around rights guarantees that operate differently from the EU's proportionality framework. In the U.S., that resistance is not legal in the direct sense because no court has blocked AI Act compliance, and the Supreme Court cannot on the analysis developed above, even if constitutional questions are being raised. It is institutional in a deeper sense, especially as the U.S. constitutional order is structured around the premise that the expressive environment of its citizens is governed by democratically accountable domestic authority. And the Brussels Effect, precisely because it operates through private market decisions abroad rather than sovereign legislative acts, cannot satisfy that premise. The result is a form of diffusion that is real in its market effects and constitutionally significant in its consequences, but democratically incomplete in its process in the U.S. Congress alone holds the authority to complete it, and the advancing compliance trajectory ensures that the political conditions for doing so will continue to intensify until it does.

VIII. Conclusion

The argument developed in this thesis resists two tempting simplifications. The first is that the Brussels Effect operates seamlessly, such that market pressure alone is sufficient to export EU regulatory norms completely and globally, and that receiving jurisdictions never play a constitutionally significant role in how those norms arrive and take hold domestically. The second is the inverse: that constitutional resistance to foreign regulatory influence renders EU diffusion structurally impossible, leaving U.S. courts as effective barriers to cross-jurisdictional norm transfer. Neither account is adequate. What this thesis has instead traced is a conditional, institutionally mediated process in which EU AI regulation encounters not a neutral surface but an active legal order, one equipped with its own doctrinal tools for assessing the conduct that foreign compliance incentives produce, and with its own institutional logic for determining where those tools fall short.

Three findings structure that account. First, the EU AI Act's core obligations, including its content prohibitions, would raise serious First Amendment difficulties if enacted by Congress. The AI Act's content-based restrictions could not survive the strict scrutiny that U.S. constitutional doctrine applies to government interference with expression, and any interpretive characterisation requirements exceed the factual and non-controversial category that *Zauderer* permits for compelled commercial disclosure. The EU legislature applied proportionality analysis calibrated to the Charter of Fundamental Rights, a framework that expressly permits balancing expression against competing rights. That analysis was neither designed to satisfy, nor tested against, the constitutional standards applicable in the U.S. What is a permissible regulatory condition within one constitutional order is a potential constitutional friction in the other.

Second, these obligations are already shaping the product architectures of AI systems serving U.S. users, with the most architecturally demanding phase of that shaping only now materialising as the Act's implementation deadlines approach. The GPAI Code of Practice engagement of leading U.S. providers, including Anthropic and OpenAI, demonstrates that the compliance trajectory is real and that the governance choices being made now are being applied globally through unified product architectures. This is the de facto Brussels Effect operating in precisely the form the theory predicts. Private commercial decisions, made in response to

market access conditions set by a foreign regulator, produce regulatory effects in jurisdictions whose governments have not acted and whose courts have not authorised those effects.

Third, and most importantly, no judicial remedy is available to address the constitutional tension that results. Because compliance with the AI Act is voluntary, the state action requirement that anchors First Amendment claims is systematically absent. As the causal chain between EU market access conditions and the expressive effects felt by U.S. users runs through the autonomous commercial judgments of private firms, the traceability requirement for a legal standing cannot be satisfied. And because the systemic dynamic of global standardisation cannot be reversed by any order a federal court could issue against any identifiable defendant, redressability fails as well. The constitutional tension is therefore real, engaging the core values the First Amendment exists to protect. Judicial remedy is structurally unavailable. What courts cannot resolve does not disappear but accumulates as pressure in the legislative domain, where the U.S. constitutional order does retain the institutional capacity to act.

That institutional logic is what makes the case for congressional action not merely available but, as the compliance trajectory advances, necessary. The democratic legitimacy deficit at the heart of the problem is precise: the content parameters shaping AI systems used by U.S. users are being calibrated to EU constitutional choices made by a legislature with no democratic accountability to U.S. citizens, applied through private compliance decisions that no U.S. institution has assessed against First Amendment standards. The First Amendment reflects a foundational commitment to the proposition that the boundaries of permissible expression in the U.S. are set by U.S. democratic deliberation and reviewable by U.S. courts. When those boundaries are instead set by a foreign legislature and diffused through private market decisions, that accountability structure is bypassed. Congress is the only institution positioned to restore it, by bringing the categories of obligation that produce First Amendment friction into the domestic legislative process where they can be tested against U.S. constitutional standards and subjected to democratic deliberation. The advancing compliance trajectory, particularly the approaching August 2026 deadline for high-risk AI obligations under Articles 9 through 15, ensures that the political conditions for doing so will continue to intensify.

These findings carry a broader lesson for how transnational regulatory diffusion is theorised. Global AI governance is not settled by the decisions of regulators alone. It is continuously renegotiated by the firms that implement compliance decisions, the courts that assess their domestic legality, and the legislatures that may ultimately intervene to resolve

tensions that litigation cannot reach. The Brussels Effect captures one powerful mechanism in that renegotiation: market access conditions, set by a jurisdiction large enough to shape global commercial behaviour, diffused through the rational compliance choices of firms that cannot afford to exit. This thesis has tried to show that the mechanism is real but not frictionless, and that the friction it encounters in constitutional receiving orders is itself analytically significant rather than incidental noise to be explained away. The non-divisibility that makes global standardisation economically rational does not make it constitutionally seamless. The constitutional orders through which EU regulatory influence must travel are active institutional environments, not passive terrain, and their response to foreign regulatory diffusion, whether through litigation, through legislative intervention, or through the structural absence of either, is part of what determines how the world comes to be governed by AI regulation that no single state designed for the world.

Bibliography

Primary Sources:

Legal Texts

- Charter of Fundamental Rights of the European Union [2012] OJ C 326/391
- 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) [2016] OJ L 119/1
 - o Cited as: GDPR
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence [2024] OJ L 1689/1
 - o Cited as: AI Act
- United States Constitution

Cases

- Allen v Wright 468 US 737 (1984)
- Blum v Yaretsky 457 US 991 (1982)
- Burton v Wilmington Parking Authority 365 US 715 (1961)
- Grayned v City of Rockford 408 US 104 (1972)
- Lujan v Defenders of Wildlife 504 US 555 (1992)
- Moody v NetChoice LLC 603 US (2024) No 22–277
- National Institute of Family and Life Advocates v Becerra 585 US (2018)
- Reed v Town of Gilbert 576 US 155 (2015)
- Simon v Eastern Kentucky Welfare Rights Organization 426 US 26 (1976)
- Snyder v Phelps 562 US 443 (2011)
- Sorrell v IMS Health Inc 564 US 552 (2011)
- United States v Playboy Entertainment Group Inc 529 US 803 (2000)
- Zauderer v Office of Disciplinary Counsel 471 US 626 (1985)

Executive Documents

Executive Office of
the President of the
U.S.

Removing Barriers to American Leadership in Artificial
Intelligence
[https://www.whitehouse.gov/presidential-
actions/2025/01/removing-barriers-to-american-leadership-in-
artificial-intelligence/](https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/)
23.1.2025
Last accessed: 31.05.2026

Executive Office of
the President of the
U.S.

Promoting the Export of the American AI Technology Stack
[https://www.whitehouse.gov/presidential-
actions/2025/07/promoting-the-export-of-the-american-ai-
technology-stack/](https://www.whitehouse.gov/presidential-actions/2025/07/promoting-the-export-of-the-american-ai-technology-stack/)
23.7.2025
Last accessed: 31.05.2026

Executive Office of
the President of the
U.S.

National Policy Framework for Artificial Intelligence: Legislative
Recommendations
[https://www.whitehouse.gov/wp-
content/uploads/2026/03/03.20.26-National-Policy-Framework-
for-Artificial-Intelligence-Legislative-Recommendations.pdf](https://www.whitehouse.gov/wp-content/uploads/2026/03/03.20.26-National-Policy-Framework-for-Artificial-Intelligence-Legislative-Recommendations.pdf)
March 2026
Last accessed: 31.05.2026

Parliamentary and Congressional Materials

Altman, Sam

Written Testimony before the Senate Judiciary Subcommittee
on Privacy, Technology and the Law
[https://www.judiciary.senate.gov/imo/media/doc/2023-05-
16%20-%20Bio%20&%20Testimony%20-%20Altman.pdf](https://www.judiciary.senate.gov/imo/media/doc/2023-05-16%20-%20Bio%20&%20Testimony%20-%20Altman.pdf)
16.5.2023
Last accessed: 31.05.2026

Amodei, Dario

Written Testimony before the Senate Judiciary Committee
[https://www.judiciary.senate.gov/imo/media/doc/2023-07-
26_-_testimony_-_amodei.pdf](https://www.judiciary.senate.gov/imo/media/doc/2023-07-26_-_testimony_-_amodei.pdf)
26.7.2023
Last accessed: 31.05.2026

*Montgomery,
Christina*

Written Testimony before the Senate Judiciary Subcommittee
on Privacy, Technology and the Law
<https://www.judiciary.senate.gov/imo/media/doc/2023-05-16%20-%20Testimony%20-%20Montgomery.pdf>
16.5.2023
Last accessed: 31.05.2026

*Russell, Jason
Matheny*

Written Testimony before the Senate Judiciary Committee
https://www.judiciary.senate.gov/imo/media/doc/2023-07-26_-_testimony_-_russell.pdf
26.7.2023
Last accessed: 31.05.2026

United States Senate
Committee on the
Judiciary

Oversight of AI: Rules for Artificial Intelligence
<https://www.judiciary.senate.gov/committee-activity/hearings/oversight-of-ai-rules-for-artificial-intelligence>
16.5.2023
Last accessed: 31.05.2026

US Congress

Winning the AI Race: Strengthening U.S. Capabilities in
Computing and Innovation
<https://www.congress.gov/119/chrg/CHRG-119shrg61426/CHRG-119shrg61426.pdf>
May 8 2025
Last accessed: 31.05.2026

US Constitution
Annotated

Amdt1.7.2.2 Vagueness, Statutory Language, and Free
Speech
https://constitution.congress.gov/browse/essay/amdt1-7-2-2/ALDE_00013539/
Last accessed: 31.05.2026

US Constitution
Annotated

Amdt1.7.13.6 Legislative Investigations
<https://www.law.cornell.edu/constitution-conan/amendment-1/legislative-investigations>
Last accessed: 31.05.2026

US Constitution
Annotated

Amdt1.7.14.1 Overview of Compelled Speech
https://constitution.congress.gov/browse/essay/amdt1-7-12-1/ALDE_00000769/
Last accessed: 31.05.2026

US Senate Committee on Commerce, Science, and Transportation	Winning the AI Race: Strengthening U.S. Capabilities in Computing and Innovation https://www.commerce.senate.gov/press/rep/release/winning-the-ai-race-strengthening-u-s-capabilities-in-computing-and-innovation-2025-4/ 29 April 2025 Last accessed: 31.05.2026
---	--

EU Institutional Documents

European Commission	Guidelines for providers and deployers of AI high-risk systems https://digital-strategy.ec.europa.eu/en/policies/guidelines-ai-high-risk-systems Last accessed: 31.05.2026
---------------------	---

European Commission	Regulatory Framework Proposal on Artificial Intelligence https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai Last accessed: 31.05.2026
---------------------	--

European Commission	Second Draft of the General-Purpose AI Code of Practice Published https://digital-strategy.ec.europa.eu/en/library/second-draft-general-purpose-ai-code-practice-published-written-independent-experts 19.12.2024 Last accessed: 31.05.2026
---------------------	--

European Commission	Signatory Taskforce for the GPAI Code of Practice https://digital-strategy.ec.europa.eu/en/policies/signatory-taskforce-gpai-code-practice Last accessed: 31.05.2026
---------------------	---

European Commission	The General-Purpose AI Code of Practice https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai Last accessed: 31.05.2026
---------------------	---

European Parliamentary
Research Service

Enforcement of the AI Act
<https://epthinktank.eu/2026/03/18/enforcement-of-the-ai-act/>
18.3.2026
Last accessed: 31.05.2026

European Parliamentary
Research Service

Implementation and Enforcement of the AI Act
[https://www.europarl.europa.eu/RegData/etudes/STUD/2025/778575/ECTI_STU\(2025\)778575_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2025/778575/ECTI_STU(2025)778575_EN.pdf)
October 2025
Last accessed: 31.05.2026

Secondary Sources:

*Almada, Marco and Petit,
Nicolas*

The EU AI Act: A Medley of Product Safety and Fundamental Rights?
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4308072
18.10.2023
Last accessed: 31.05.2026

*Almada, Marco and Radu,
Anca*

The Brussels Side-Effect: How the AI Act Can Reduce the Global
Reach of EU Policy
2024
Last accessed: 31.05.2026

Balkin, Jack M

Moody v NetChoice: The Supreme Court Meets the Free Speech
Triangle
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5104013
19.1.2025
Last accessed: 31.05.2026

Bradford, Anu

The Brussels Effect
<https://www.proquest.com/docview/1286682884>
2012
Last accessed: 31.05.2026

Bradford, Anu

Digital Empires: The Global Battle to Regulate Technology (Oxford
Academic 2023)

Bradford, Anu

The Brussels Effect: How the European Union Rules the World (1st
edn, Oxford University Press 2020)

- Bryson, Joanna J, Diamantis, Mihailis E and Grant, Thomas D* Liability for Artificial Intelligence in Dubber, Markus D, Pasquale, Frank and Das, Sunit (eds), The Oxford Handbook of Ethics of AI
28 October 2022
Last accessed: 31.05.2026
- Calo, Ryan* Artificial Intelligence Policy: A Primer and Roadmap' 51(2) UC Davis Law Review
<https://digitalcommons.law.uw.edu/cgi/viewcontent.cgi?article=1639&context=faculty-articles>
2017
Last accessed: 31.05.2026
- Clark, Sam, Haeck, Pieter and Carlson, Kathryn* EU Pressure Builds on Anthropic Over Mythos Hacking Risks Politico
<https://www.politico.eu/article/eu-pressure-builds-on-anthropic-over-mythos-hacking-risks/>
5.5.2026
Last accessed: 31.05.2026
- Crowell & Moring LLP* The US Senate Hearing on Artificial Intelligence and Intellectual Property
<https://www.crowell.com/en/insights/client-alerts/the-us-senate-hearing-on-artificial-intelligence-and-intellectual-property>
6.12.2023
Last accessed: 31.05.2026
- Crum, Bastiaan* Brussels Effect or Experimentalism? The EU AI Act and Global Standard-Setting 14(3) Internet Policy Review
<https://policyreview.info/pdf/policyreview-2025-3-2032.pdf>
27.8.2025
Last accessed: 31.05.2026
- Damro, Chad* Market Power Europe 19(5) Journal of European Public Policy
<https://doi.org/10.1080/13501763.2011.646779>
1.3.2012
Last accessed: 31.05.2026
- Daskal, Jennifer* Borders and Bits 71(1) Vanderbilt Law Review
<https://www.proquest.com/docview/2040740630>
January 2018
Last accessed: 31.05.2026

- Douek, Evelyn and Lakier, Genevieve* Rereading "Editorial Discretion" (Knight First Amendment Institute)
<https://knightcolumbia.org/blog/rereading-editorial-discretion>
 24 October 2022
 Last accessed: 31.05.2026
- Eurostat* Population on 1 January 2026
<https://doi.org/10.2908/TPS00001>
 Last accessed: 31.05.2026
- McDermott Will & Emery* No State AI Law Moratorium in One Big Beautiful Bill Act
<https://www.mcdermottlaw.com/insights/no-state-ai-law-moratorium-in-one-big-beautiful-bill-act/>
 5.7.2025
 Last accessed: 31.05.2026
- Metz, Cade and Conger, Kate* Anthropic Faces Scrutiny Over Claude Mythos, The New York Times
<https://www.nytimes.com/2026/05/12/technology/anthropic-claude-mythos.html>
 12.5.2026
 Last accessed: 31.05.2026
- Michelman, Frank I* The State Action Doctrine in Amar, Vikram D and Tushnet, Mark V (eds) Global Perspectives on Constitutional Law (Oxford University Press 2009)
- Pero, James* Europe Can't Get Meta Ray-Ban Display Because of EU Regulations' Gizmodo
<https://gizmodo.com/privacy-isnt-the-only-thing-plaguing-metas-smart-glasses-2000737911>
 25.3.2026
 Last accessed: 31.05.2026
- Machado, Priscilla H* State Action and the Supreme Court EBSCO Research Starters
<https://www.ebsco.com/research-starters/law/state-action-and-supreme-court>
 2023
 Last accessed: 31.05.2026
- Ropes & Gray LLP* Examining the Landscape and Limitations of the Federal Push to Override State AI Regulation
<https://www.ropesgray.com/en/insights/alerts/2026/03/examining-the-landscape-and-limitations-of-the-federal-push-to-override-state-ai-regulation>
 11.3.2026
 Last accessed: 31.05.2026

- Rubinfeld, Jed* The First Amendment's Purpose 53 Stanford Law Review
<https://heinonline-org.uaccess.univie.ac.at/HOL/Page?collection=journals&handle=hein.journals/stflr53&id=786>
 4. 2001
 Last accessed: 31.05.2026
- Schauer, Frederick* The Exceptional First Amendment' KSG Working Paper No RWP05-021
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=668543
 18.2.2005
 Last accessed: 31.05.2026
- Schauer, Frederick* Towards an Institutional First Amendment 89 Minnesota Law Review
<https://scholarship.law.umn.edu/cgi/viewcontent.cgi?article=1681&context=mlr>
 2005
 Last accessed: 31.05.2026
- Schwartz, Paul M* Global Data Privacy: The EU Way' (2019) 94(4) New York University Law Review
<https://nyulawreview.org/wp-content/uploads/2019/10/NYULAWREVIEW-94-4-Schwartz.pdf>
 2019
 Last accessed: 31.05.2026
- Siegmann, Charlotte and Anderljung, Markus* The Brussels Effect and Artificial Intelligence: How EU Regulation Will Impact the Global AI Market (Governance of AI Research Paper)
https://cdn.governance.ai/Brussels_Effect_GovAI.pdf
 16.8.2022
 Last accessed: 31.05.2026
- Tsesis, Alexander* Compelled Speech and Proportionality 97(3) Indiana Law Journal art 1
<https://www.repository.law.indiana.edu/ilj/vol97/iss3/1>
 2022
 Last accessed: 31.05.2026
- US Department of Commerce Data Privacy Framework Program Overview
<https://www.dataprivacyframework.gov/Program-Overview>
 Last accessed: 31.05.2026

Vaassen, Bram

AI, Opacity, and Personal Autonomy 35 Philosophy and Technology

<https://doi.org/10.1007/s13347-022-00577-5>

22.9.2022

Last accessed: 31.05.2026