POLICY STUDY

The European Union's
ARTIFICIAL INTELLIGENCE ACT

AN EXTRATERRITORIAL INCURSION ON AMERICANS' INALIENABLE RIGHTS

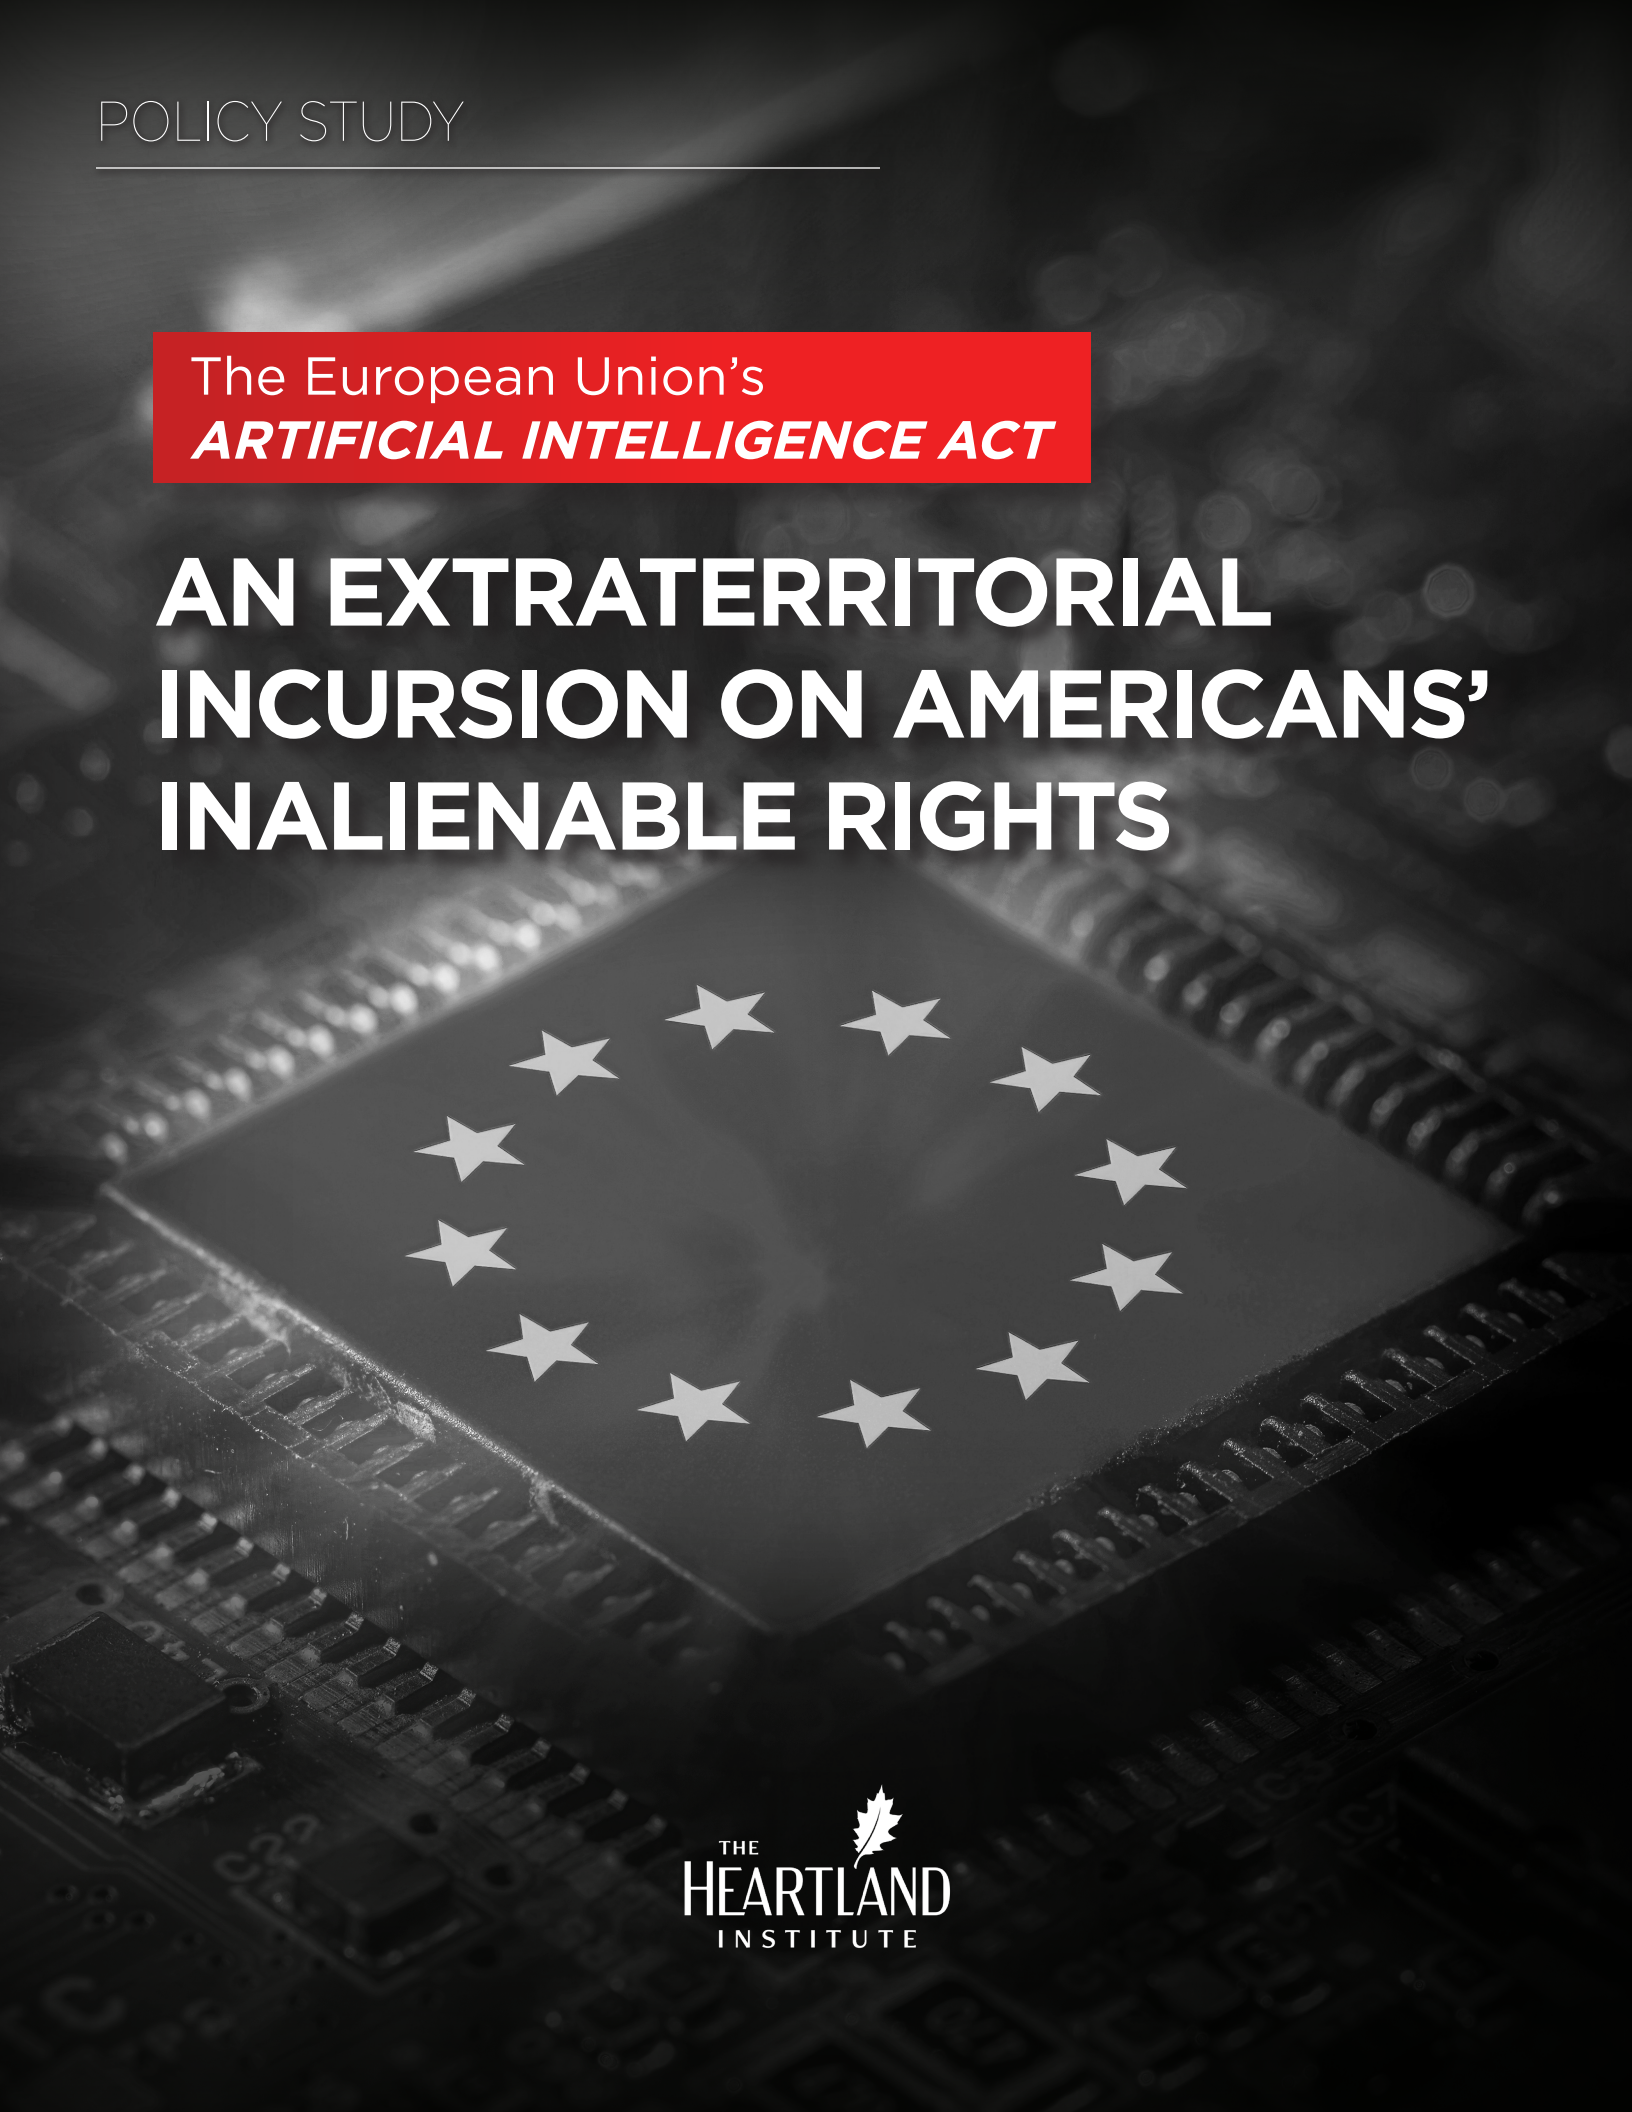

THE HEARTLAND INSTITUTE

Jack McPherrin is a research fellow within The Heartland Institute's Glenn C. Haskins Emerging Issues Center (EIC), and also serves as Heartland's research editor, managing the editorial process for all internally and externally authored Heartland research and policy publications. As an EIC research fellow, Jack authors research papers, policy studies, legislative tip sheets, opinion and analysis articles, and other research and policy publications focusing on a variety of emerging topics that pose threats to America's fundamental principles, values, and institutions.

For more information, please call 312/377-4000 or visit our website at Heartland.org.

1933 N. Meacham Road, Suite 550, Schaumburg, IL, 60173

POLICY STUDY

The European Union's
ARTIFICIAL INTELLIGENCE ACT

AN EXTRATERRITORIAL INCURSION ON AMERICANS' INALIENABLE RIGHTS

CONTENTS

Key Takeaways

- The European Union's Artificial Intelligence Act establishes a centralized, extraterritorial, global regulatory model for AI. It applies whenever outputs are "used in the Union," allowing the EU to impose its standards globally through turnover-based penalties and value chain obligations.
- The law divides AI into risk tiers—unacceptable (banned), high-risk (heavily regulated), limited-risk (transparency only), and minimal-risk—plus a separate regime for general-purpose AI (GPAI) covering foundation models.
- Prohibited uses include behavioral manipulation and distortion, exploitation of vulnerabilities, perpetuating discrimination, social scoring, predictive policing by profiling, large-scale facial-recognition databases, and emotion inference in workplaces or schools.
- High-risk systems—covering areas such as biometrics, education, employment, essential services, law enforcement, and elections—must undergo risk management, data-bias mitigation, human oversight, conformity assessments, CE-marking, and ongoing monitoring. Deployers must often complete a Fundamental Rights Impact Assessment (FRIA) before use.
- GPAI providers must publish training-data summaries, adopt EU-compliant copyright rules, and prepare technical documentation covering how the model was built and evaluated.
- GPAIs deemed to pose "systemic risk," which incorporates practically all frontier large language models (LLMs) currently on the marketplace, must run adversarial tests, continuously assess and mitigate risks, report serious incidents to the EU AI Office, and ensure robust cybersecurity protections.
- Penalties reach up to 7 percent of global turnover for violating prohibited practices and 3 percent of global turnover for most other violations. Blanket prohibitions are already in effect; GPAI and high-risk obligations phase in at specified intervals from 2025–2027.
- Implications for the United States:
    - Sovereignty: EU regulators effectively set global AI standards without U.S. representation.
    - Speech and Expression: "Manipulation," "exploitation," and "bias" provisions pressure platforms to suppress controversial expression.
    - Competition and Innovation: Compliance costs entrench large incumbents and deter startups and open research, while slowing development and innovation.
- Policy priorities: The White House, Congress, and the states should (1) leverage trade negotiations to reject foreign regulatory reach over U.S. firms, (2) craft legislation targeting truly coercive uses without restricting inalienable rights and preempting cooperation with the EU, and (3) advance innovation-driven, voluntary standards rooted in American principles.

Introduction

Europe's New Digital Dominion

The European Union's Artificial Intelligence Act (AI Act) represents a watershed moment in the global regulation of emerging technologies.[1] Enacted in 2024 and already entering into force, it is the world's first comprehensive legal framework governing the development and use of artificial intelligence. Yet beneath the rhetoric of ethics, safety, and "trustworthy AI," the AI Act imposes something far more consequential: regulatory authority that extends far beyond Europe's borders.

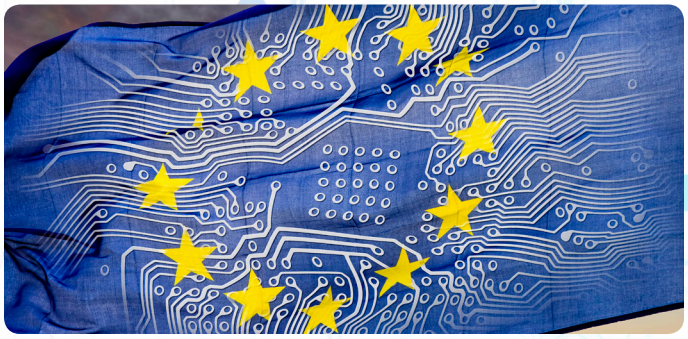

The European Union's Artificial Intelligence Act (AI Act) represents a watershed moment in the global regulation of emerging technologies. Enacted in 2024 and already entering into force, it is the world's first comprehensive legal framework governing the development and use of artificial intelligence. Yet beneath the rhetoric of ethics, safety, and "trustworthy AI," the AI Act imposes something far more consequential: regulatory authority that extends far beyond Europe's borders.

By claiming jurisdiction over entities whose AI systems are used within the European Union (EU), Brussels has effectively arrogated the power to dictate how Americans design, deploy, and interact with AI technologies. The law thus embodies a growing European tendency to export its technocratic vision to the rest of the world, subordinating national sovereignty and individual liberty to the diktats of the EU bureaucracy. This pattern is neither new nor confined to the realm of artificial intelligence. The EU has already attempted to impose extraterritorial control through measures such as the Corporate Sustainability Reporting Directive (CSRD) and the Corporate Sustainability Due Diligence Directive (CSDDD), which require nonEuropean firms to conform to the EU's environmental, social, and governance (ESG) mandates—backed by the

1 Except where otherwise noted, all references in this paper refer to the European Union's Artificial Intelligence Act and are cited by article and/or section number (e.g., "EU AI Act, Article 2, Paragraph 1"). See: European Union, Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (*Artificial Intelligence Act*), *Official Journal of the European Union,* L 2024/1689, July 12, 2024, https://eur-lex.europa.eu/eli/reg/2024/1689/oj

threat of crippling penalties.[2]

The AI Act follows the same playbook. It claims to protect consumers and uphold "fundamental rights," but in practice, it extends European governance into the global digital economy by redefining what constitutes acceptable technology and permissible speech. Under its expansive, riskbased framework, the EU positions itself as the moral arbiter of algorithmic behavior, wielding broad discretion to determine which applications of AI are "manipulative," "biased," or "untrustworthy." Such power, exercised by unelected officials across the Atlantic, stands in direct contradiction to American constitutional principles of free expression, due process, and limited government.

This paper provides an overview of the AI Act's central provisions—explaining how the law classifies and regulates AI systems under the guise of risk management, and outlining the specific obligations imposed on companies, including penalties for noncompliance—before turning to a broader analysis of what this framework means for the United States, particularly for the preservation of free speech, innovation, and sovereignty in the digital age.

This paper provides an overview of the AI Act's central provisions—explaining how the law classifies and regulates AI systems under the guise of risk management, and outlining the specific obligations imposed on companies, including penalties for noncompliance—before turning to a broader analysis of what this framework means for the United States, particularly for the preservation of free speech, innovation, and sovereignty in the digital age.

Many of the Act's most consequential provisions have already taken effect, with further obligations rolling out over the next two years. As the European Union moves forward with its enforcement, understanding the scope and implications of this law is critical to safeguarding Americans' rights and ensuring foreign bureaucrats do not dictate the terms of our digital future.

2 For an in-depth analysis of the European Union's Corporate Sustainability Due Diligence Directive (CSDDD) and its implications for the United States, see: Jack McPherrin and Justin Haskins, "CSDDD: The European Union's Corporate Sustainability Due Diligence Directive Is a Direct Threat to U.S. Sovereignty, Free Markets, and Individual Liberty," *Policy Study,* The Heartland Institute, March 31, 2025, https://heartland.org/publications/csddd-the-european-unions-corporate-sustainability-due-diligence-directive-is-a-direct-threat-to-u-s-sovereignty-free-markets-and-individual-liberty/

Scope and Structure

Who and What the AI Act Covers

The EU's Artificial Intelligence Act casts a far-reaching regulatory net. It applies not only to European developers and deployers of AI systems but also to any entity—based inside or outside the European Union—whose AI products are placed on the EU market or whose outputs are used within the Union. Specifically, the law covers providers, deployers, importers, distributors, manufacturers, and authorized representatives of providers—regardless of whether these entities are located in the European Union. It also applies to "affected persons" located within the European Union.[3]

In other words, the law's reach extends to American firms whose AI tools or models are accessed by European users, integrated into European services, or generate outputs consumed by EU individuals or institutions. This extraterritorial clause ensures that European regulators may claim jurisdiction over companies that have no physical presence in Europe and no intent to target European consumers, so long as their systems exert an effect within the bloc's borders.

This structure mirrors the EU's broader approach to global governance: using access to its single market to impose European standards on the world. The implications for American companies are clear. U.S. firms could be compelled to abide by European definitions of discrimination, manipulation, or bias—concepts that are not only vague but deeply political. An AI-generated recommendation protected as free expression

OpenAI CEO Sam Altman

PHOTO COURTESY WORLD ECONOMIC FORUM/FLICKR.COM

The law's reach extends to American firms whose AI tools or models are accessed by European users, integrated into European services, or generate outputs consumed by EU individuals or institutions. This extraterritorial clause ensures that European regulators may claim jurisdiction over companies that have no physical presence in Europe and no intent to target European consumers, so long as their systems exert an effect within the bloc's borders.

3 EU AI Act, Article 2(1).

under the First Amendment, for instance, could be deemed "manipulative" or "exploitative" under EU standards.[4] Such asymmetry is heightened by the law's exemptions for military, defense, and national-security uses, as well as for certain law-enforcement activities,[5] illustrating that Brussels has carved out privileges for its own governmental apparatus while extending its authority abroad.

Once an entity falls within the Act's scope, its obligations depend on how its AI systems are classified under a tiered, risk-based framework. This structure lies at the heart of the regulation and determines both the degree of oversight and the compliance burdens imposed. The Act essentially addresses four main levels of risk: unacceptable, high, limited, and minimal, while also covering general-purpose AI (GPAI) models in a similar risk-based approach.

Once an entity falls within the Act's scope, its obligations depend on how its AI systems are classified under a tiered, risk-based framework. This structure lies at the heart of the regulation and determines both the degree of oversight and the compliance burdens imposed. The Act essentially addresses four main levels of risk: unacceptable, high, limited, and minimal, while also covering general-purpose AI (GPAI) models in a similar risk-based approach.

1. Unacceptable Risk

Systems deemed to present an "unacceptable risk" are prohibited outright.[6] These include applications that:

- Deploy subliminal or manipulative techniques with the objective or the effect of "materially distorting the behaviour of a person or a group of persons by appreciably impairing their ability to make an informed decision."
- Exploit vulnerabilities of individuals or groups of individuals "due to their age, disability or a specific social or economic situation."
- Evaluate or classify individuals or groups of individuals based on social scores.
- Predict the occurrence of an individual committing a crime based solely on profiling or assessing personal characteristics.
- Create or expand facial recognition databases through scraping images from the internet or CCTV footage.
- Infer individuals' emotions in the workplace or educational settings.
- Use biometric categorization to deduce or infer race, political opinions, beliefs, sexual orientation, or other characteristics, or use biometric identification for law enforcement, with narrow exceptions.

Though some of these prohibitions are clearly beneficial, the breadth of others gives regulators significant interpretative discretion, such as what constitutes exploiting vulnerabilities or manipulating groups of people.

4 EU AI Act, Article 5(1).

5 EU AI Act, Article 2(3) and Article 2(4).

6 EU AI Act, Article 5.

2. High Risk

Systems classified as "high risk" are the primary focus of the AI Act and are subject to the Act's most stringent obligations, which will be discussed in the next section. This category covers two main tracks: (1) AI systems that serve as safety components in regulated products such as medical devices or vehicles,[7] and (2) systems deployed in sensitive domains.[8] Enumerated in Annex III, these sensitive domains include:

- Biometrics.
- Critical infrastructure.
- Education and vocational training.
- Employment, workers' management, and access to self-employment.
- Access to "essential private services and essential public services and benefits," such as health care services, credit, insurance, and first responders.
- Law enforcement.
- Migration, asylum, and border control.
- Administration of justice and democratic processes.

The European Commission retains the authority to amend the annexes, meaning the already-broad list of "high-risk" uses can expand over time.[9]

3. Limited Risk

A third group of AI systems does not trigger the full high-risk regime but still raises concerns about opacity. For those systems, the Act imposes transparency obligations only:

- Providers must design those AI systems that interact directly with people in a way that users are told they are dealing with AI, unless it is already obvious.[10]
- Providers must ensure that AI systems that generate synthetic or manipulated audio, image, video, or text clearly disclose that the output is artificial.[11]
- Deployers using emotion-recognition or biometric-categorization tools must inform affected persons.[12]
- Deployers publishing deepfakes or other manipulated media must disclose that fact.[13]

While the Act does not use the term "limited risk," this category is widely described as such in the Commission and industry guidance. It applies to chatbots, generative content tools, and other applications that could potentially mislead users without explicit disclosure.

7 EU AI Act, Article 6(1).

8 EU AI Act, Article 6(2) and Annex III.

9 EU AI Act, Article 7.

10 EU AI Act, Article 50(1).

11 EU AI Act, Article 50(2).

12 EU AI Act, Article 50(3).

13 EU AI Act, Article 50(4).

4. Minimal Risk

While the Act does not explicitly identify a fourth category, it encourages the residual AI systems—those that fall outside the three tiers mentioned above—to voluntarily adopt comparable best practices. These systems are effectively unregulated, a gap the Commission expects to fill through voluntary codes of conduct and ethical guidelines.[14] They are likely to include ordinary consumer and enterprise tools such as spam filters, productivity software, and video-game algorithms that do not process sensitive data or produce legally significant outcomes.

General-Purpose AI

Finally, the Act introduces a separate framework for general-purpose AI (GPAI) models, the large, multipurpose systems that form the backbone of today's AI ecosystem.[15] These include well-known models such as ChatGPT, Claude, and Gemini, which can perform a wide range of tasks and be adapted for countless applications. Because GPAI models are not designed for any single use, the Act regulates them differently from the risk tiers above.

As will be discussed more in the next section, the AI Act establishes baseline duties for all GPAI providers—such as transparency through technical documentation and training data summaries—and additional oversight for the largest or most powerful models deemed to present "systemic risk." GPAI does not constitute a new risk tier. Rather, it is a cross-cutting category that can underpin any of the four levels described above. A GPAI model might power a high-risk medical-diagnostic tool, a transparency-only chatbot, or a minimal-risk consumer app—each of which would still be regulated according to its own classification under the risk framework.

14 EU AI Act, Recital 165.

15 EU AI Act, Articles 51-56.

Obligations Under the EU AI Act

As noted in the previous section, limited- and minimal-risk AI systems are largely subject only to disclosure obligations. The highest-risk category—AI systems that perpetuate "prohibited practices"—has no compliance strictures because these uses are simply banned. As such, the teeth of the AI Act bite in two places: (1) the high-risk regime that governs design, documentation, testing, oversight, and post-market control across the entire AI value chain; and (2) a separate framework for GPAI models, especially those designated as posing "systemic risk." Together, these rules effectively export EU processes, values, and standards onto the global market, including U.S. companies—even if such companies have no physical presence in the European Union.

These rules effectively export EU processes, values, and standards onto the global market, including U.S. companies—even if such companies have no physical presence in the European Union.

High-Risk AI Systems

Before placement on the EU market or use in the European Union, every high-risk system must meet several basic requirements:

- **Risk-management:** The company must identify foreseeable harms to safety and fundamental rights, plan how to reduce them, and keep that plan up to date.[16]
- **Data governance:** Training and testing data must be suitable for the stated purpose and handled in ways that reduce errors and unfair bias.[17]
- **Technical documentation:** The system needs a clear technical file—paperwork that explains how it was designed, tested, and will be overseen—so regulators can check it for compliance.[18]
- **Record-keeping:** Systems must keep logs so important decisions can be reconstructed if something goes wrong.[19]

16 EU AI Act, Article 9.

17 EU AI Act, Article 10.

18 EU AI Act, Article 11.

19 EU AI Act, Article 12.

- **Instructions for use:** The company must provide instructions that explain what the system can and cannot do and how to use it safely.[20]
- **Human oversight:** The system's design must allow people to supervise it and step in when needed.[21]
- **Accuracy, robustness, and cybersecurity:** Performance and security must be "fit for purpose."[22]

Beyond these baseline requirements, the Act delineates responsibilities by role, depending on where an entity falls in the AI system's value chain—from creation to real-world use and application. Four specific roles are identified: the provider (the company that offers the system under its name), the importer (who brings it into the EU), the distributor (who sells or resells it inside the EU), and the deployer (the organization that uses it in practice). A U.S. developer selling directly to an EU customer will usually be the provider; if it sells through an EU partner, that partner may be the importer or distributor. Either way, obligations are not limited to the developer. They are spread across the value chain.

For providers of high-risk systems, the law requires a structured, up-front check called a "conformity assessment" to show the product meets the Act's baseline requirements: risk management, suitable data practices, documentation, human oversight, accuracy, robustness, and security. When the provider has completed those checks, it issues an EU "declaration of conformity" and places the CE mark— best understood as the EU's "passport" that signals the system passed the required review—on the product. After launch, providers must monitor real-world performance, correct problems, and report serious incidents to national authorities on set timelines.[23]

Beyond these baseline requirements, the Act delineates responsibilities by role, depending on where an entity falls in the AI system's value chain—from creation to real-world use and application. Four specific roles are identified: the provider (the company that offers the system under its name), the importer (who brings it into the EU), the distributor (who sells or resells it inside the EU), and the deployer (the organization that uses it in practice).

Importers and distributors act as gatekeepers inside Europe. Importers must verify that the system they are bringing in is compliant by ensuring that the CE mark and documents exist, and that storage or transport won't undermine compliance.[24] Distributors must refrain from placing non-compliant systems on the market and are expected to pause sales and escalate concerns if something looks wrong.[25] In practice, this means EU partners will ask U.S. providers for straightforward proof that the pre-market checks were done and that clear instructions exist for safe use.

20 EU AI Act, Article 13.

21 EU AI Act, Article 14.

22 EU AI Act, Article 15.

23 EU AI Act, Article 16.

24 EU AI Act, Article 23.

25 EU AI Act, Article 24.

The deployer—the bank, hospital, school, agency, or firm that turns the system on—also has concrete duties. Deployers must use the system as instructed, ensure trained human oversight, keep relevant logs so important decisions can be reconstructed, and check the quality of input data where it materially affects outcomes.[26] For certain high-risk uses—especially by public bodies or providers of public-facing services—the deployer must also complete a Fundamental Rights Impact Assessment (FRIA) before the system goes live. A FRIA is a comprehensive pre-deployment analysis that explains what the system will do, who could be affected, and how it could impact affected persons' rights.[27]

In short, the high-risk regime is not a single checklist but a layered set of pre-market, contractual, and ongoing operational duties that bind every actor from developer to end-user. For U.S. companies whose tools are used in the European Union, this means front-loaded compliance reviews, proof-of-compliance demands from EU partners, customer requests for FRIA inputs, and continuing obligations to monitor, fix, and report serious incidents.

General-Purpose AI Systems

Separate from the high-risk regime, the Act creates baseline duties for developers of general-purpose AI (GPAI) models—the large, flexible models that can be adapted for many tasks. All GPAI providers must:

- Prepare technical documentation that explains how the model was built and evaluated.
- Give downstream developers enough information and instructions to use the model responsibly.
- Adopt a policy to respect EU copyright rules.
- Publish a "sufficiently detailed" summary of the model's training-data sources.[28]

These obligations, however, become substantially more demanding if a GPAI model is deemed to pose "systemic risk." The Act defines systemic risk as the danger that a model's high-impact capabilities could cause widespread harm to the Union's market, or cause "actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or the society as a whole," particularly when such effects can "propagate at scale" through downstream applications.[29]

The Act's recitals illustrate the kinds of harms captured under "systemic risk," which include, but are not limited to:

- Negative effects in relation to major accidents.
- Disruptions of critical sectors and serious consequences to public health and safety.
- Actual or reasonably foreseeable negative effects on democratic processes.
- Actual or reasonably foreseeable negative effects on public and economic security.
- The dissemination of illegal, false, or discriminatory content.[30]

This definition is notably broad, effectively granting the Commission wide latitude to determine what constitutes "systemic" harm—which, in practice, captures nearly every leading general-purpose model on the market today.

26 EU AI Act, Article 26.

27 EU AI Act, Article 27.

28 EU AI Act, Article 53.

29 EU AI Act, Article 3(65).

30 EU AI Act, Recital 110.

A model is classified as entailing systemic risk under two circumstances: if a model has either (a) "high impact capabilities," defined by the law as being above a computational threshold of 10^{25} floating-point operations (FLOPs) or (b) if the European Commission designates the model as posing systemic risk on its own initiative or based on the recommendation of a scientific panel.[31]

Once a GPAI model falls into this category, its provider must:

- Conduct standardized evaluations and document adversarial ("red-team") tests.
- Continuously assess and mitigate systemic risks at the EU level.
- Report serious incidents to the EU AI Office and national authorities "without undue delay."
- Ensure robust cybersecurity protections.
- Demonstrate compliance through harmonized standards or interim codes of practice until such standards are adopted.[32]

The AI Act therefore operates on two intersecting tracks: a system-level regime for high-risk uses and a model-level regime for GPAI. For U.S. firms whose products or models reach EU users, both tracks impose recurring obligations—at contract signature, before deployment, and after launch. Those workloads translate into ongoing documentation, testing, and data-handling duties in agreements with customers and vendors. Failure to meet these duties exposes firms to substantial, revenue-linked penalties, as detailed in the next section.

The AI Act therefore operates on two intersecting tracks: a system-level regime for high-risk uses and a model-level regime for GPAI. For U.S. firms whose products or models reach EU users, both tracks impose recurring obligations—at contract signature, before deployment, and after launch. Those workloads translate into ongoing documentation, testing, and data-handling duties in agreements with customers and vendors. Failure to meet these duties exposes firms to substantial, revenue-linked penalties, as detailed in the next section.

31 EU AI Act, Article 51.

32 EU AI Act, Article 55.

Enforcement and Penalties

The enforcement provisions of the AI Act are as far-reaching as its substantive obligations. While the European Commission sets overall policy direction, day-to-day enforcement rests primarily with national regulators of individual EU member states known as "national competent authorities," including at least one "notifying authority" and at least one "market surveillance authority."[33] These national regulators are responsible for monitoring compliance, conducting investigations, and imposing penalties for most violations. They are also tasked with maintaining a central EU database where providers of high-risk systems must register their products before they are placed on the market.[34]

At the supranational level, the European Commission—through its newly created EU AI Office—oversees enforcement related to GPAI models, particularly those designated as posing systemic risk. The Commission may request documentation, evaluate models directly, and coordinate with national authorities to ensure consistency across the Union.[35]

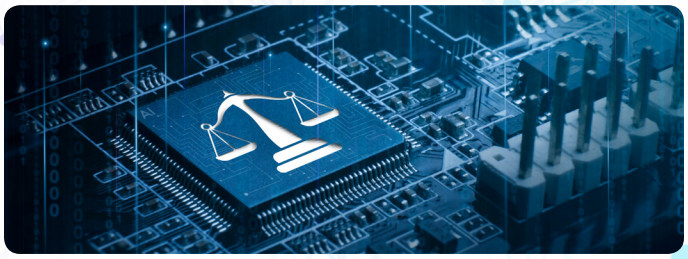

These rules effectively export EU processes, values, and standards onto the global market, including U.S. companies—even if such companies have no physical presence in the European Union.

Penalty Structure

The AI Act introduces a tiered system of monetary penalties that mirror—and in some cases exceed—the structure of the EU's other extraterritorial regulations, such as the Corporate Sustainability Due Diligence Directive. The fines are pegged to a company's total worldwide annual turnover—similar to a company's total revenue—ensuring that non-European companies face enormous liability if their systems are used within the EU.

The maximum penalties are as follows:

- **Prohibited AI practices:** up to €35 million or 7 percent of global annual turnover, whichever is higher.[36]

33 EU AI Act, Article 70.

34 EU AI Act, Article 71.

35 EU AI Act, Articles 64-66.

36 EU AI Act, Article 99(3).

- **Violations of obligations under the Act (high-risk and GPAI)**: up to €15 million or 3 percent of global annual turnover, whichever is higher.[37]
- **Supplying incorrect, incomplete, or misleading information to regulators or notified bodies**: up to €7.5 million or 1 percent of global annual turnover, whichever is higher.[38]
- **For small and medium-sized enterprises (SMEs)**, the same percentage thresholds apply, but authorities apply the lower total.[39]

Beyond fines, authorities can order corrective actions, restrict or withdraw non-compliant systems, and require disclosures, which can be as costly as the penalties themselves. Notably, these penalties can be triggered by user-generated content. Providers are responsible for ensuring that their systems do not produce outputs deemed to engage in prohibited practices, even when those outputs originate from independent users inside the European Union.

Extraterritorial Reach

Perhaps most strikingly, these penalties apply extraterritorially. Any company—European or otherwise—whose AI outputs are used within the EU can be investigated, fined, or compelled to take remedial action. The Act's jurisdictional provisions rely not on a company's geographic location but on whether its system is "used in the Union."[40] This means that U.S. companies can face fines of up to 7 percent of global turnover for actions taken entirely on U.S. soil, provided their models or products generate outputs accessed by European users. For example, if an American user of a GPAI model creates content deemed to be "exploitative" or "manipulative" that later circulates widely online—including among EU residents—the model's provider could be held liable under the Act and be fined up to 7 percent of its global turnover. More detailed examples of how this could work in practice will be provided later in this paper.

In short, the AI Act establishes a multi-tiered enforcement regime that combines national and supranational regulators backed by substantial turnover-based fines. Together, these mechanisms give Brussels and individual EU countries both the institutional capacity and financial leverage to police AI systems worldwide. For American developers and service providers, the practical effect is clear: even without a European presence, they are subject to a foreign compliance regime that blends technocratic oversight with the threat of crippling pecuniary sanctions.

The AI Act establishes a multi-tiered enforcement regime that combines national and supranational regulators backed by substantial turnover-based fines. Together, these mechanisms give Brussels and individual EU countries both the institutional capacity and financial leverage to police AI systems worldwide. For American developers and service providers, the practical effect is clear: even without a European presence, they are subject to a foreign compliance regime that blends technocratic oversight with the threat of crippling pecuniary sanctions.

37 EU AI Act, Article 99(4).

38 EU AI Act, Article 99(5).

39 EU AI Act, Article 99(6).

40 EU AI Act, Article 2.

Implementation and Enforcement Timeline

Though the AI Act entered into force on August 1, 2024, its obligations apply in stages.[41] This phased rollout gives the European Commission and national regulators time to build the machinery of enforcement and gradually expand the law's reach. For American developers and firms whose systems are used in Europe, these milestones show how the AI Act's obligations expand year by year, gradually pulling more companies, sectors, and models into its scope.

The first phase began six months after entry into force, on February 2, 2025, when prohibitions on AI systems deemed to pose "unacceptable risk" took effect. As noted in earlier sections of this paper, this includes bans on manipulative or deceptive systems, social-scoring applications, and certain forms of biometric surveillance.[42]

The second phase arrived one year after entry into force, on August 2, 2025, when obligations for GPAI systems and related governance provisions began to apply. This is also the date on which the European Commission's AI Office assumed supervision of GPAI models, and when individual EU countries must have established their market-surveillance frameworks. It also marks the first date from which penalties can be assessed against companies for violating the prohibited practices introduced in phase one.[43]

The third phase will arrive two years after entry into force, on August 2, 2026, when the bulk of the

> **The law's reach extends to American firms whose AI tools or models are accessed by European users, integrated into European services, or generate outputs consumed by EU individuals or institutions. This extraterritorial clause ensures that European regulators may claim jurisdiction over companies that have no physical presence in Europe and no intent to target European consumers, so long as their systems exert an effect within the bloc's borders.**

41 EU AI Act, Article 113.

42 EU AI Act, Article 113, Article 4.

43 EU AI Act, Article 113.

Act's remaining provisions become enforceable. This phase includes the beginning obligations and associated penalties for "high-risk" systems delineated by Annex III—namely, the sectoral high-risk uses such as hiring, credit, education, health care, law enforcement, migration, critical infrastructure, and democratic processes. This is also the date on which penalties against non-compliant GPAI companies become enforceable.[44]

The fourth and final phase will arrive three years after entry into force, by August 2, 2027. This phase requires that the remaining high-risk systems—those used as a safety component for products like medical devices and vehicles—and all GPAI models placed on the market before August 2025 be brought into full compliance.[45]

Though additional deadlines extend into the late 2020s and early 2030s—including provisions related to evaluation, reporting, and large-scale public-sector systems[46]—these four phases capture the core implementation milestones that define the Act's practical impact.

The core timeline can thus be summarized as follows:

- **February 2025** (6 months after entry into force): Blanket prohibitions take effect.
- **August 2025** (12 months): GPAI obligations; penalties for prohibited practices take effect.
- **August 2026** (24 months): Obligations and penalties for most high-risk systems; penalties for GPAI systems take effect.
- **August 2027** (36 months): Obligations and penalties for remaining high-risk systems and legacy GPAI models take effect.

These deadlines illustrate how rapidly the EU is moving from enactment to enforcement. Prohibitions and transparency rules have already begun to shape the behavior of global AI developers. By 2026-2027, nearly every significant AI system operating in or affecting Europe will fall within the law's regulatory orbit. Understanding these stages is essential to grasping the full scope of the Act's extraterritorial reach.

Together, these deadlines illustrate how rapidly the EU is moving from enactment to enforcement. Prohibitions and transparency rules have already begun to shape the behavior of global AI developers. By 2026-2027, nearly every significant AI system operating in or affecting Europe will fall within the law's regulatory orbit. Understanding these stages is essential to grasping the full scope of the Act's extraterritorial reach.

44 EU AI Act, Article 113.

45 EU AI Act, Article 113, Article 111(3).

46 EU AI Act, Article 111.

Implications and Analysis

The central question for the United States is not whether Europe may regulate American-made AI products and services sold and used in Europe, but whether the EU can shape how American companies design and operate AI products used everywhere else, even in the United States. Because the Act's jurisdictional trigger is use "in the Union" rather than corporate domicile, European standards apply whenever outputs cross the Atlantic.[47] Coupled with turnover-based penalties and shared obligations across the value chain, that structure pushes global providers to align to the strictest interpretation and then apply those constraints worldwide. This dynamic, familiar from Europe's prior forays into global rulemaking, carries profound implications for American sovereignty, speech, competition, and innovation.

The European Union's attempt to institutionalize global control over artificial intelligence follows a pattern of regulatory extraterritoriality that is neither novel nor mysterious. In fact, this pattern—known as the "Brussels Effect"—is well-documented by many reputable academic sources. Moreover, this strategy has escalated in recent years, with extraterritorial EU laws such as the Corporate Sustainability Reporting Directive (CSRD), Corporate Sustainability Due Diligence Directive (CSDDD), and General Data Protection Regulation (GDPR) being passed and entering into various phases of force.

Sovereignty and Rulemaking Without Representation

The European Union's attempt to institutionalize global control over artificial intelligence follows a pattern of regulatory extraterritoriality that is neither novel nor mysterious. In fact, this pattern—known as the "Brussels Effect"—is well-documented by many reputable academic sources.[48] Moreover, this strategy has escalated in recent years, with extraterritorial EU laws such as the Corporate Sustainability Reporting Directive (CSRD), Corporate Sustainability Due Diligence Directive

47 EU AI Act, Article 2.

48 See Anu Bradford, *The Brussels Effect: How the European Union Rules the World*, 2020, New York, NY: Oxford University Press, https://scholarship.law.columbia.edu/books/232/

(CSDDD),[49] and General Data Protection Regulation (GDPR) being passed and entering into various phases of force. The GDPR's extraterritorial scope pushed global firms to adopt EU-style privacy defaults worldwide,[50] while the CSRD/CSDDD regime extends EU environmental, social, and governance reporting and due-diligence requirements deep into non-EU supply chains. The AI Act hews to the same script.

Importantly, global companies have a strong and rational incentive to align with the EU standard by default. Some may think that the solution would simply be for companies to run separate "EU" and "U.S." versions that are compliant with each jurisdiction's regulations, but this rarely works in practice. Two sets of rules, settings, and contracts are costly to maintain and difficult to keep perfectly separated. One mistake—an EU user accessing the "U.S." model or a reseller in Europe distributing a non-EU product—would trigger substantial penalties. To avoid this, most companies would simply standardize to the strictest rule set, because it is the most conservative, risk-averse strategy.

The result is a quiet transfer of authority over key judgments—what AI may generate, when persuasion becomes "manipulation," which inferences are "unacceptable"—from American institutions to European regulators and notified bodies. No U.S. legislature has made these choices, yet they begin to bind American firms the moment their outputs are "used in the Union." In day-to-day business, EU customers and partners will also demand contractual proof of compliance throughout supply and distribution chains, extending Brussels' leverage into American product design and vendor terms.

That is the sovereignty problem in plain terms: rules that reshape U.S. products and speech norms are being set by institutions entirely unaccountable to American voters and guided by substantially different legal and cultural priorities. With that foundation set, the most immediate effect Americans are likely to experience is infringement upon their rights to speech and expression.

That is the sovereignty problem in plain terms: rules that reshape U.S. products and speech norms are being set by institutions entirely unaccountable to American voters and guided by substantially different legal and cultural priorities. With that foundation set, the most immediate effect Americans are likely to experience is infringement upon their rights to speech and expression.

A De Facto Global Speech Regime

Although the AI Act is framed as product safety, its practical effect is to regulate speech by proxy. Providers and deployers are expected to prevent or mitigate outputs labeled "manipulative," "exploitative," or "biased," and can face exposure even if the content originates with independent users. Faced with turnover-based fines and coordinated oversight, global platforms will rationally calibrate models and policies to the strictest plausible European interpretation. That is the "Brussels Effect," now applied to AI-mediated expression.

49 Jack McPherrin and Justin Haskins, "CSDDD: The European Union's Corporate Sustainability Due Diligence Directive Is a Direct Threat to U.S. Sovereignty, Free Markets, and Individual Liberty."

50 European Union, "What is GDPR, the EU's new data protection law?" accessed October 25, 2025, https://gdpr.eu/what-is-gdpr/

Two scenarios illustrate how this could impact Americans.

Example 1: Climate and Energy Debate Constrained by 'Sustainability' Risk Rules

A global general-purpose AI system such as ChatGPT, integrated into search engines, office software, and content platforms, fields millions of prompts each day about climate, energy, and environmental policy. After the AI Act takes effect, European authorities issue guidance that misinformation or "harmful" outputs concerning sustainability goals may constitute manipulation of user behavior or exploitation of vulnerabilities—particularly when models influence consumption or investment decisions.

In response, the provider's compliance teams classify climate-related discussion as a "risk domain." To demonstrate alignment with EU law, they retrain moderation and ranking systems to block or heavily qualify responses that question net-zero feasibility, the cost of renewable mandates, or the efficacy of carbon-sequestration measures unless the model cites approved international sources. Because the same model weights serve both European and American users, these global filters now govern what U.S. and other non-EU policymakers, journalists, and researchers see when they query the model about climate or energy policy. Outputs increasingly mirror EU regulatory framing, presenting contested policy judgments as scientific consensus.

No European official has censored Americans directly. The chilling effect arises from the provider's attempt to minimize exposure under the Act's open-ended "risk-mitigation" duties. The combination of use-based jurisdiction and revenue-linked penalties ensures that the most restrictive European interpretation becomes the global speech baseline. What began as a European compliance exercise has become a global content filter governing how the world discusses energy policy.

Example 2: Election Commentary Throttled During an EU Campaign

A U.S. startup offers a writing assistant that generates political op-eds and social media posts. The tool is marketed for general political commentary and is accessible in the European Union. During a European Parliament campaign, users in Vienna prompt it to draft posts alleging corruption among a slate of candidates. EU authorities receive complaints that the system's outputs influence voter behavior during an election period.

Under Annex III of the AI Act, AI systems "intended to be used for influencing the outcome of an election or referendum or the voting behaviour of natural persons in the exercise of their vote in elections or referenda" are classified as high-risk. Because the model's design and marketing foresee political uses, and its foreseeable use includes persuasive messaging, the national market-surveillance authority determines that it falls within that category. The firm must now register the system in the EU database, perform a conformity assessment, implement risk-management and human-oversight measures, maintain event logs, and file documentation on data quality and model performance. Failure to comply carries penalties that would likely put the company out of business.

To satisfy these obligations, the provider adds new safety layers that restrict or flag any content capable of "influencing electoral outcomes." Those filters are built into the core model and deployed globally, since maintaining a separate U.S. system would be prohibitively expensive. As a result, American users now encounter blocked or qualified outputs whenever the system deems a prompt politically persuasive—whether it concerns Congress, state elections, or ballot initiatives.

The pressure also travels through contracts. U.S. publishers, campaign consultants, and nonprofits using the tool receive new terms of service prohibiting "use in political advocacy" and requiring audit logs of "election-related interactions." The

provider's engineers, compliance teams, and lawyers enforce European definitions of "democratic integrity" across all markets.

Again, no European official has censored Americans directly. Yet through high-risk classification and value-chain obligations, European regulators have effectively dictated how an American political-speech technology tool must behave domestically and worldwide. What began as an EU safeguard for local elections ends as a de facto global code of permissible political expression.

Two features of the law drive these outcomes. First, the jurisdictional trigger is "use in the Union,"[51] not corporate domicile. Second, obligations sit upstream with providers of both high-risk systems and general-purpose AI; high-risk systems must satisfy risk-management, data governance, technical documentation, logging, transparency, and human-oversight requirements,[52] while GPAI models face documentation, evaluations and adversarial testing, incident reporting, and "appropriate" risk-mitigation duties—with additional oversight where models are deemed to pose systemic risk.[53] When combined with prohibitions on manipulative or exploitative systems and EU-level coordination through the AI Office and national market-surveillance authorities,[54] the path of least resistance is broad, pre-emptive filtering. In short, to keep EU access and reduce regulatory ambiguity, platforms will narrow what their models will say—everywhere.

For Americans, the consequence is straightforward. AI-mediated communication of every kind—political commentary and advocacy, journalism, art, education, humor, personal advice, or everyday conversation—will be shaped by global providers aligning to European standards, even when those forms of expression are constitutionally protected in the United States. What gets suppressed is not limited to unlawful deception or targeted harassment; it includes controversial, value-laden, or unconventional ideas that might be alleged to "distort behavior," "exploit vulnerabilities," or "perpetuate discrimination" under expansive interpretations.

Over time, this could narrow the range of permissible thought and expression available to anyone who relies on mainstream AI tools. Rather than a direct European ban on U.S. speech, the Act's mechanism is subtler: risk-averse platform governance induced by extraterritorial obligations and revenue-linked penalties, which gradually export European speech norms to the rest of the world.

Over time, this could narrow the range of permissible thought and expression available to anyone who relies on mainstream AI tools. Rather than a direct European ban on U.S. speech, the Act's mechanism is subtler: risk-averse platform governance induced by extraterritorial obligations and revenue-linked penalties, which gradually export European speech norms to the rest of the world.

Market Consolidation and Cost to Innovation

The AI Act's complexity and compliance costs will not fall evenly. In practice, they favor the largest and most politically connected companies: those

51 EU AI Act, Article 2.

52 EU AI Act, Articles 8-27.

53 EU AI Act, Articles 51-56.

54 EU AI Act, Articles 64-71.

best equipped to absorb paperwork, legal review, and continuous oversight. Meeting the Act's requirements will demand entire departments for documentation, risk management, and conformity assessments. Large firms can treat this as the price of doing business. Smaller developers cannot.

For startups, the choice is bleak: scale back features, delay releases, or block European users altogether. Some will close their models to the public to control downstream use; others will abandon promising ideas before reaching market. The result is fewer competitors and greater concentration of power among firms that can afford to comply. These same firms will also shape the technical standards and guidance documents that define future obligations, further entrenching their dominance under the banner of "safety."

Innovation slows under this kind of regulatory pressure. Entrepreneurs spend time managing audits instead of building products. Investment shifts toward firms that already have compliance infrastructure. Open research communities—where breakthroughs often occur—face growing uncertainty about liability. Over time, the incentive to experiment shrinks, and the next generation of innovators must seek permission before they can take risks.

Innovation slows under this kind of regulatory pressure. Entrepreneurs spend time managing audits instead of building products. Investment shifts toward firms that already have compliance infrastructure. Open research communities—where breakthroughs often occur—face growing uncertainty about liability. Over time, the incentive to experiment shrinks, and the next generation of innovators must seek permission before they can take risks.

This outcome should concern policymakers as much as any direct speech restriction. When regulation rewards size and punishes openness, it transforms a competitive marketplace into a managed industry that is shaped by bureaucratic politics rather than organic discovery. The United States, whose comparative advantage has long been innovation through competition, should be especially alert to that risk.

Concluding Recommendations

Not every aspect of the AI Act is misguided. Some restrictions—such as limits on clearly coercive, deceptive, or exploitative applications—address legitimate dangers and reflect concerns any responsible government should take seriously. Transparency and accountability in high-risk settings have their place. The problem is that the European Union has extended those principles far beyond genuine safety, constructing a system of pre-approval, continuous oversight, and open-ended interpretation. What could have been targeted safeguards has become a mechanism for centralized control—one that the United States must understand and resist if it is to preserve self-government, individual rights, innovation, and prosperity.

To defend those interests, the United States should act on three fronts:

1. **Through trade negotiations** and related instruments, the White House and Congress should make clear that the United States will not accept the extraterritorial application of EU law to American citizens, firms, or products.

2. **Congress should adopt federal legislation** that explicitly prohibits foreign regulatory enforcement against U.S. entities operating solely within domestic borders and preempts state or corporate cooperation with such enforcement.

3. **States should follow Texas' lead** by enacting laws that safeguard Americans' constitutional rights while establishing narrow, clearly defined guardrails against genuinely harmful AI practices.[55] Broad state-level adoption of such measures would lay the groundwork for a coherent federal framework rooted in American values of liberty, innovation, and accountability.

Together, these steps would reassert American sovereignty in the digital sphere and ensure that the future of artificial intelligence is governed by the principles of free expression, competitive enterprise, and democratic accountability—not by bureaucratic decree from Brussels.

Together, these steps would reassert American sovereignty in the digital sphere and ensure that the future of artificial intelligence is governed by the principles of free expression, competitive enterprise, and democratic accountability—not by bureaucratic decree from Brussels.

55 A detailed discussion of the Texas Responsible Artificial Intelligence Governance Act (HB 149, 2025) is outside the scope of this paper, though it will appear in a forthcoming Heartland Institute publication. In the meantime, the full text of the law is available at: https://capitol.texas.gov/BillLookup/Text.aspx?LegSess=89R&Bill=HB149

THE HEARTLAND INSTITUTE

3939 North Wilke Road
Arlington Heights, IL 60004

Heartland.org

For more information on this topic, visit Heartland.org, email Think@heartland.org, or call (312) 377-4000.