

AN AMERICAN’S GUIDE TO THE EU AI ACT

Margot E. Kaminski[†] & Andrew D. Selbst^{‡‡}

The EU AI Act entered into force in August 2024. The AI Act is long. It is complicated. It relies on a regulatory framework and institutions unfamiliar to many in the United States. But as the first omnibus AI regulation worldwide, it has the potential to have a vast influence on both practice and lawmaking.

In this Article, we provide the American’s Guide to the EU AI Act. This Article breaks down the AI Act for a U.S. law audience, explaining the overall mechanisms, and how the Act interacts with background EU laws and institutions. At its core, the AI Act is structured on Europe’s product safety regime. It is aimed at governing AI systems through assigning them into risk tiers, and deploying bans, risk regulation, and self-regulation. But it also contains later-drafted provisions on general-purpose AI that depart from this framework, as well as multiple ad-hoc provisions and other regulatory strands.

The Article also analyzes the consequences of framing AI regulation as risk regulation and of constructing AI systems as products rather than bureaucratic processes. It describes the AI Act itself as a “legal exoskeleton,” with hard law built around the softer belly of technical standards. The Article identifies the threats this poses for both substance and legitimacy, and the potential political ramifications of that design.

TABLE OF CONTENTS

I.	INTRODUCTION	1082
II.	SOME NECESSARY BACKGROUND ON EU INSTITUTIONS, EU LAW, AND MEMBER STATE INSTITUTIONS.....	1083
A.	EU INSTITUTIONS	1083
B.	EU LAWS	1087
C.	MEMBER STATE INSTITUTIONS	1091
III.	THE LAW	1092
A.	THE CORE PRODUCT SAFETY REGIME: THE “NEW LEGISLATIVE FRAMEWORK”	1093
1.	<i>The Risk Tiers</i>	1094
2.	<i>The Bans</i>	1095
3.	<i>High-Risk AI Systems</i>	1098

DOI: <https://doi.org/10.15779/Z38JS9HB12>

© 2025 Margot E. Kaminski and Andrew D. Selbst.

[†] Moses Lasky Professor of Law, Colorado Law School, and Director of the Privacy Initiative at Silicon Flatirons Center. Recipient of a 2024 Fulbright-Schuman grant and Fernand Braudel Senior Fellowship at the European University Institute (EUI). Thanks to Deirdre Curtin for hosting and involving Professor Kaminski in the EUI intellectual community and related discussions. Thanks to Marco Almada and Nicolas Petit for detailed and thoughtful feedback on this piece. Mistakes are all ours.

^{‡‡} Professor of Law, University of California, Los Angeles, School of Law.

a)	What Practices Are “High Risk”?	1099
b)	What Are the Substantive Requirements for Providers (Developers) of High-Risk AI Systems?	1101
c)	What Are The Substantive Requirements for Deployers (Users) of High-Risk AI Systems?	1105
d)	Conformity Assessments	1106
e)	Accountability and Enforcement.....	1108
B.	GENERAL-PURPOSE AI MODELS	1111
C.	AD HOC ELEMENTS AND OTHER STRANDS OF REGULATION	1115
IV.	ANALYSIS.....	1118
A.	PRECAUTION, OR A BID FOR AI BUSINESS?	1119
B.	LEGALLY CONSTRUCTING HIGH-RISK AI SYSTEMS THROUGH PRODUCT SAFETY	1120
C.	THE ACT AS PRODUCT OF ITS DRAFTING STORY	1126
D.	THE ACT AS LEGAL EXOSKELETON.....	1126
E.	IN WHICH IT ALL COMES DOWN TO POWER POLITICS.....	1132
V.	CONCLUSION	1133

I. INTRODUCTION

The EU AI Act entered into force in August 2024. The AI Act is long. It is complicated. It relies on a regulatory framework and institutions unfamiliar to many in the United States. But as the first omnibus AI regulation worldwide, it has the potential to have a vast influence on both practice and lawmaking.

In this Article, we provide an American’s Guide to the EU AI Act. We begin in Part II with necessary background on EU law and institutions that readers can skim or skip, if they are already familiar with the European Union (EU). In Part III, we give an overview of the AI Act. At its core, the AI Act is structured on Europe’s product safety regime. It is aimed at governing predictive AI systems through assigning them into risk tiers, and deploying bans, risk regulation, and self-regulation. But it also contains later-drafted provisions on general-purpose AI that depart from this framework, as well as multiple ad hoc provisions.

In Part IV, we offer our analysis. We point to the consequences of framing AI regulation as risk regulation and of constructing AI systems as products rather than bureaucratic processes. We describe the AI Act itself as a “legal exoskeleton,” with hard law built around the softer belly of technical standards. We identify the threats this poses for both substance and legitimacy, and the potential political ramifications of that design.

We close in Part V with a word of warning about global power politics. This may not be the time for another legal export out of Brussels. And the AI Act itself is not well-designed to be exported. Instead, we caution that deregulatory forces are compounding, both within Europe and from the United States.

II. SOME NECESSARY BACKGROUND ON EU INSTITUTIONS, EU LAW, AND MEMBER STATE INSTITUTIONS

The AI Act sits atop a mountain of existing EU law and institutional structure. In this section, we begin with some basic and not-so-basic background that we think is necessary to understand what's going on with the AI Act.

A. EU INSTITUTIONS

We start with EU institutions, briefly covering both the basics and how these institutions are relevant to the AI Act. The EU is a supranational organization consisting of a system of twenty-seven member states, joined in a common legal and economic enterprise. An American audience might understand it as a kind of federalist system, though perhaps one more akin to the one imagined by the Articles of Confederation, with member states exercising true sovereignty.

Certain lawmaking institutions operate at the EU level by either directly legislating or delegating to member states. The EU has four primary decision-making bodies: the European Council, the European Parliament, the Council of the European Union, and the European Commission. The European Council—not to be confused with the Council of the EU¹—comprises the EU heads of state. It sets the general political direction of the EU but is not generally involved in legislation. The two main legislative bodies of the EU are the European Parliament and the Council of the European Union. The Parliament represents the citizens of member states and is directly elected by them. The Council of the EU consists of ministers from each member state and differs depending on the policy area of the law being considered. The primary executive body of the EU is the European Commission.

We offer the following overview of the EU legislative process because the AI Act in several places envisions bypassing it for purposes of amendments or

1. It's also important not to confuse either with the "Council of Europe," a totally separate international institution comprising 46 countries, including all EU members, and dedicated to upholding human rights and democracy. We know this is a big ask, because wow, that's a lot of different things named the "Council." But, you know, just try.

implementation. As a general matter, new EU laws are typically passed through what is known as the “ordinary legislative process,” in which the Commission uses its “right of initiative” to propose a new law, which is then taken up jointly by the Parliament and Council of the EU. A new law will often undergo a lengthy “trilogue” negotiation, an informal institutional negotiation in which members of the three bodies come together to work out a draft, which can then be formally adopted by each of the three bodies internally. This is not unlike a bill going to conference in the U.S. Congress, except more complicated, as there are three bodies that need to agree. This is the process that the AI Act followed in its three-year drafting history.

The Commission (executive) has certain specific roles designated by the AI Act, including the delegated ability to modify certain aspects of the law without going through the full legislative process.² The Act also created an “AI Office” within the Commission to “develop Union expertise and capabilities in the field of AI.”³ This European AI Office has already been involved, primarily as a convener, in drawing up a General-Purpose AI Code of Practice.⁴ The AI Office has also released a template for summarizing training data for general-purpose AI.⁵

Back to EU law: the powers of the EU are conferred by treaty. Unlike EU member states whose sovereignty is assumed, the European Union does not have inherent powers, because it is not a sovereign state.⁶ The EU’s power to

2. Regulation (EU) 2024/1689 of the European Parliament and of the Council, art. 97, 2024 O.J. (L 2024/1689) [hereinafter AI Act].

3. *Id.* art. 64.

4. *Drawing-Up a General-Purpose AI Code of Practice*, EUR. COMM’N (Aug. 1, 2025), <https://digital-strategy.ec.europa.eu/en/policies/ai-code-practice> (“The AI Office played a pivotal role throughout the process . . . facilitating the drawing-up, coordinating the discussions and documenting the outcomes”). On August 1, 2025, the “Commission and AI Board approve[d] the code via Adequacy Decisions.”

5. *See* EUROPEAN AI OFFICE, EUROPEAN AI OFFICE WORKING GROUP MEETINGS, CODE OF PRACTICE FOR GENERAL PURPOSE AI: TEMPLATE FOR SUMMARY OF TRAINING DATA (Jan. 17, 2025), <https://ec.europa.eu/newsroom/dae/redirection/document/111909>; *see also* *Drawing-Up AI Code of Practice*, *supra* note 4 (“[T]he AI Office is also developing a template on the sufficiently detailed summary of training data that general-purpose AI model providers are required to make public according to Article 53(1)d) of the AI Act . . . [T]he AI Office has presented its preliminary ideas and allowed the participants to the Code to provide additional feedback on the preliminary structure and elements of the template . . . [and] was also discussed with the Member States’ representatives in the AI Board subgroup and the European Parliament before the Commission adopts the template in the second quarter of 2025.”).

6. This principle is referred to as the “principle of conferral.” *See* Consolidated Version of the Treaty on European Union art. 5(2), June 7, 2016, 2016 O.J. (C 202) 13 [hereinafter TFEU] (“Under the principle of conferral, the Union shall act only within the limits of the competences conferred upon it by the Member States in the Treaties to attain the objectives

make law—referred to as its “competences”—varies in different subject matter areas, depending on what treaties confer to it.⁷ In some areas, the EU’s power is exclusive, while in others it is shared with or supporting of the member states.⁸ For example, the EU has exclusive power over the Euro and governing trade; shared power over consumer protection, some types of social policy, and governance of the internal market more generally; and supporting power only for areas such as industry, culture, and tourism.⁹ Member states retain power for internal security and crime prevention,¹⁰ and in general, any power not explicitly conferred on the EU belongs to member states.¹¹

While this “principle of conferral” formally still governs, it does seem that the EU is undergoing an expansion of authority similar to the U.S. federal government’s under the Commerce Clause.¹² Many different legal issues are considered regulation of the “internal market,” and thus can be subject to EU power.¹³ The AI Act, as explained below, is part of a larger EU product safety regime justified as a regulation of the internal market.¹⁴

set out therein. Competences not conferred upon the Union in the Treaties remain with the Member States.”).

7. *See id.* title I, art. 2.

8. *Id.* art. 2(1) (“When the Treaties confer on the Union exclusive competence in a specific area, only the Union may legislate and adopt legally binding acts, the Member States being able to do so themselves only if so empowered by the Union or for the implementation of Union acts.”); *id.* art. 2(2) (“When the Treaties confer on the Union a competence shared with the Member States in a specific area, the Union and the Member States may legislate and adopt legally binding acts in that area.”); *id.* art. 2(5) (“In certain areas and under the conditions laid down in the Treaties, the Union shall have competence to carry out actions to support, coordinate or supplement the actions of the Member States, without thereby superseding their competence in these areas.”).

9. *See generally infra* Part III; *see also* TFEU, *supra* note 6.

10. *See* TFEU, *supra* note 6, art. 72 (“This Title shall not affect the exercise of the responsibilities incumbent upon Member States with regard to the maintenance of law and order and the safeguarding of internal security.”); *see also* TFEU, *supra* note 6, arts. 275–76 (describing the limits of the jurisdiction of the European Court of Justice).

11. *See id.* arts. 4(1), 5(2); Sacha Garben, *Competence Creep Revisited*, 57 J. COMMON MKT. STUD. 205, 205 (Mar. 2019) (“[P]owers that have not been explicitly conferred on the EU remain exclusively with the Member States (Article 4(1) and 5(2) TFEU).”).

12. *See, e.g.*, ROBERT SCHUTZE, INTRODUCTION TO EUROPEAN LAW 64 (4th ed. 2023) (“[T]hree developments have led to widespread accusations that the European Union’s competences are ‘unlimited.’”).

13. Garben, *supra* note 11, at 207 (“The Treaty’s functional powers—mostly, but not exclusively related to the internal market—can cut horizontally through all policy areas, including those where the EU has no, or only complementary, competence. This means that the EU can, through such indirect powers, legislate in areas that are considered to fall within national autonomy.”).

14. *See* AI Act, *supra* note 2, art. 1 (describing the purpose of the regulation as, among other things, “to improve the function of the internal market.”).

The principal EU court is the Court of Justice of the European Union (CJEU).¹⁵ The CJEU works in parallel with national courts of member states to effectuate EU law.¹⁶ One of its main functions is to interpret EU law. The CJEU is one institution but has two court with different jurisdiction: the European Court of Justice (ECJ) and the General Court. The ECJ primarily hears cases referred to them by the high courts of member states, including preliminary rulings on new issues of law. The General Court, by contrast, can hear directly from individuals and institutions, but only for claims that an EU body directly violated their rights.

Generally, if an individual wants to raise a claim that private parties or national institutions have violated EU law, she must go through her national court.¹⁷ National courts hear disputes on the application of EU law to their citizens and can refer an issue of European law to the ECJ. This is referred to as the “preliminary reference procedure.”¹⁸ The role of the ECJ can vary depending on the structure of a particular treaty. The ECJ has an active role, for example, in data protection law. Under the General Data Protection Regulation,¹⁹ a data subject can file a complaint with the local data protection authority, and if dissatisfied, can sue that authority in national court with the ECJ as a potential backstop. But even under the GDPR, which grants data subjects many individual rights, a data subject may not sue in the ECJ directly.

By contrast, although the AI Act is purportedly motivated in large part by concern for individual fundamental rights, there is no individual right of

15. Not to be confused with the European Court of Human Rights, the court that oversees the European Convention on Human Rights, a treaty passed by the Council of Europe. *See generally* ROBERT SCHUTZE, EUROPEAN UNION LAW, ch. 10: Judicial Powers I: (Centralized) European Procedures (3d ed. 2021).

16. SCHUTZE, *supra* note 15, ch. 10 (“In addition to a number of direct actions (direct actions start directly in the European Court), the EU Treaties here envisage an indirect action starting in the national courts: the preliminary reference procedure. This procedure is the judicial cornerstone of the Union’s cooperative federalism. For it combines the central interpretation of Union law by the Court of Justice with the decentralised application of European law by the national courts.”).

17. Individuals may bring an “action for annulment” of EU law under art. 263 of TFEU, asking for a particular EU law to be annulled. *See* Action for Annulment, EU: Summaries of EU Legislation (Oct. 29, 2010), <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=legisum:ai0038>. However, the Court favors the preliminary reference procedure. *See* SCHUTZE, *supra* note 15, at 373. And if an individual wants to challenge a purported violation of EU law, they use preliminary reference.

18. *See* TFEU, *supra* note 6, art. 267(1)(b); *see also* Preliminary Ruling Proceedings—Recommendations to National Courts, EU: Summaries of EU Legislation (Dec. 3, 2024), <https://eur-lex.europa.eu/EN/legal-content/summary/preliminary-ruling-proceedings-recommendations-to-national-courts.html>.

19. Regulation (EU) 2016/679 of the European Parliament and of the Council, 2016 O.J. (L 119) 1 [hereinafter GDPR].

redress in the Act. The institutions involved are focused on the operation of products in the market rather than hearing cases on fundamental rights. Thus, the role of the CJEU in performing oversight or backstopping much of the law is less clear.²⁰

We have thus far covered the main lawmaking, executive, and judicial institutions of the EU. The AI Act also refers to several other institutions that may be unfamiliar to a U.S. audience. The AI Act relies on the European Data Protection Supervisor (EDPS), a body that exists to enforce data protection law, to serve a few different functions. The European Data Protection Supervisor is designated as the overseer where an EU-level body acts in a capacity that would normally be overseen by a member state institution such as a market surveillance authority.²¹ It is also designated as an observer for the EU AI Board created by the law.²²

Then there are the European Standards Organizations (ESOs). These are private organizations, not governmental, but they are considered European bodies and are governed by EU law. There are three standards bodies authorized by certain EU laws to create technical standards that help effectuate the AI Act. Each has a slightly different area of expertise. While the European Committee for Standardization (CEN) is more generalized, the European Committee for Electrotechnical Standardization (CENELEC) focuses on standards related to electrical engineering, and the European Telecommunications Standards Institute (ETSI) focuses on information and communications.²³ The AI Act calls for and heavily relies on “technical” standards development, and has enlisted a joint technical committee of CEN and CENELEC to implement the AI Act’s standards.²⁴

B. EU LAWS

Here, we turn from EU institutions to EU laws. We first cover the meaning of commonly used terms, like regulations, directives, and recitals. Then we turn to the broader substantive legal setting behind the AI Act, identifying several relevant and overlapping EU laws.

EU statutory laws come in two principal forms: regulations and directives. A regulation is a law that is directly binding on individual entities of the EU—

20. Some of the provisions of the AI Act may be found to have direct effect and thus be subject to being raised before national courts. Thanks to Nicholas Petit for this important caveat.

21. AI Act, *supra* note 2, art. 74(9).

22. *Id.* art. 65.

23. *European Standardization*, CEN-CENELEC, <https://www.cencenelec.eu/european-standardization/cen-and-cenelec/>.

24. *Artificial Intelligence*, CEN-CENELEC, <https://www.cencenelec.eu/areas-of-work/cen-cenelec-topics/artificial-intelligence/>.

citizens, companies, and member states. A directive is, by contrast, a law that directs member states to pass laws on a topic, within certain parameters. If you think of the EU as supranational, then regulations are akin to self-executing treaties, while directives are akin to non-self-executing treaties. A principal function of EU law in general is the harmonization of member state laws. A directive is used when the EU feels that a lighter touch on harmonization is required, whereas a regulation is passed to achieve tighter harmonization.

For example, the GDPR was passed in 2016 because the data protection laws of member states under the 1995 Data Protection Directive (DPD)²⁵ were widely fragmented as implemented.²⁶ Also, a directive will sometimes serve as a stepping stone to a later regulation if deemed necessary by the legislative bodies, as in the transition from the DPD to the GDPR, or in the case of the e-Commerce Directive²⁷ and the Digital Services Act.²⁸

EU laws are made up of articles, recitals, and sometimes annexes. The articles make up the binding legislative text. Recitals, by contrast, illustrate the purpose behind the law; they are like formalized legislative history. But unlike legislative history in the United States—which is often denigrated as either a last resort or irrelevant to interpretation—recitals are vitally important. The CJEU operates under a theory of purposive interpretation, in which text is to be interpreted in light of its declared purpose.²⁹ Thus, recitals play a direct role in giving meaning to EU law.

Some EU laws contain annexes, which are also binding legislative text, written separately. An article may refer to an annex for a procedure or list of covered circumstances to apply a particular provision. The separation allows for easier updates of implementing details. Annexes are heavily used in the AI Act.

Now we turn from structure to legal substance. The AI Act implicitly and explicitly relies on, overlaps with, and is constrained by other EU laws. The most relevant of these is a model known as the “New Legislative Framework” (NLF)—the framework for the EU’s product safety regime. Laws built on the

25. Directive 95/46/EC, of the European Parliament and of the Council, 1995 O.J. (L 281) 31.

26. GDPR, *supra* note 19, recital 9 (“[O]bjectives and principles of Directive 95/46/EC remain sound, but it has not prevented fragmentation in the implementation of data protection across the Union.”).

27. Directive 2000/31/EC, of the European Parliament and of the Council (Directive on electronic commerce), 2000 O.J. (L 178) 1.

28. Regulation (EU) 2022/2065, of the European Parliament and of the Council (Digital Services Act), 2022 O.J. (L 277) 1.

29. See, e.g., Gerard Conway, *The Quality of Decision-Making at the Court of Justice of the European Union*, in *HOW TO MEASURE THE QUALITY OF JUDICIAL REASONING* 225, 227 (Mátyás Bencze & Gar Yein Ng eds., 2018).

NLF operate by requiring products to meet certain “essential requirements” for safety before entering the market.³⁰ They are allowed to enter the market only after a “conformity assessment” is performed, showing that the product—as well as its risk mitigation framework, failure alert procedures, and documentation—meet the framework’s essential requirements. Typical NLF laws regulate children’s toys³¹ or elevators³²—standard product safety regimes that might be assigned to the Consumer Product Safety Commission in the United States. Familiarity with the NLF is principally important because the AI Act was created as a product safety law under the NLF and uses the same conformity assessment and market surveillance oversight structures that the rest of the NLF laws use.³³ We go into more specifics below when we discuss the substance of the law.

Other laws are also important to the AI Act for different reasons. The Charter of Fundamental Rights of the European Union lays out the EU’s fundamental rights regime.³⁴ The AI Act relies on the Charter and interpreting cases implicitly in two ways. First, the Act indicates that AI poses risks to fundamental rights, but does not delineate which rights are affected or how they might be affected by AI. It therefore leans on the existing EU fundamental rights framework to fill this gap. The Act also does not offer individual rights of redress, instead deferring to other existing laws to enable individual redress for violations of fundamental rights.

The data protection regime of the EU, specifically the GDPR, is a particularly important complement to the AI Act. While AI is built on data, the AI Act is not a data protection law. For that, there is the GDPR. Data protection is a fundamental right protected by the Charter.³⁵ The GDPR is backstopped by the CJEU, which interprets the GDPR’s regulatory

30. Stéphane du Boispi  an, Markus Mueck & Christophe Gaie, *Introduction to the European New Legislative Framework*, in EUROPEAN DIGITAL REGULATIONS 1, 2–3 (Markus Mueck & Christophe Gaie eds., 2025).

31. Directive 2009/48/EC of the European Parliament and of the Council (on the safety of toys), 2009 O.J. (L 170) 1.

32. Directive 2014/33/EU of the European Parliament and of the Council, 2014 O.J. (L 96) 251.

33. The NLF has also recently been applied to “products with digital elements” through the Cyber Resilience Act (Regulation (EU) 2024/2847 of the European Parliament and of the Council, 2024 O.J. (L. Series)). According to Marco Almada, in communications with the Authors, “[i]t does not face many of the issues created by the AI Act, as its regulatory object is much more technical.” (communication on file with authors). But the Cyber Resilience Act, too, raises concerns about using a risk-based approach to fundamental rights. See Pier Giorgio Chiara, *Understanding the Regulatory Approach of the Cyber Resilience Act: Protection of Fundamental Rights in Disguise?*, 16 EUR. J. RISK REGUL. 469 (2025).

34. Charter of Fundamental Rights of the European Union, 2000 O.J. (C 364) 1.

35. *Id.* arts. 7, 8.

requirements in light of fundamental rights. Thus, even where the CJEU may appear to be absent from the AI Act, it is actively present in a closely complementary and overlapping law.

The GDPR and AI Act each seek to mitigate or protect against different types of harms to different sets of people: harms to data subjects versus harms to those affected by AI. However, some aspects of the laws directly overlap, for example where personal data is used to make decisions about individuals. Some requirements like the right to explanation of an AI decision compete or overlap.³⁶

The AI Act also creates exceptions to some GDPR rules. For example, it allows the processing of “special category” data under the GDPR (i.e., sensitive data like race) “when strictly necessary for the purpose of ensuring bias detection and correction,” as long as certain safeguards are followed.³⁷ The Act also permits the processing of personal data for regulatory sandboxing—essentially AI pilot programs with safe harbors.³⁸ Finally, as mentioned above, the Act designates roles for the European Data Protection Supervisor—an office created by the GDPR—including making the EDPS the relevant oversight authority for EU entities that would be regulated by the law.³⁹

The AI Act also interacts substantially with the Digital Services Act. The DSA covers the use of platforms, and those that use AI for content moderation will be subject both to the DSA and the AI Act. In that case, the DSA is the *lex generalis* to the AI Act’s *lex specialis*, and the platform AI will be governed by the DSA, with the AI Act’s requirements riding atop it. The same principle governs AI that happens to also be a product covered by an existing NLF regime, like a toy or a drone. The AI Act instructs providers to comply with the existing NLF framework, but to add on the AI Act’s requirements.⁴⁰ Finally, the Act also overlaps with EU copyright law, especially the Copyright in the Digital Single Market Directive.⁴¹

36. Compare Case C-203/22, *CK v Magistrat der Stadt Wien*, ECLI:EU:C:2025:117 (Feb. 27, 2025), with AI Act, *supra* note 2, art. 86; see also Margot E. Kaminski & Gianclaudio Malgieri, *The Right to Explanation in the AI Act*, in *THE EU ARTIFICIAL INTELLIGENCE ACT: A THEMATIC COMMENTARY* (Gianclaudio Malgieri, Gloria González Fuster, Alessandro Mantelero & Gabriela Zanfir-Fortuna eds., forthcoming 2026) (describing the “hydraulic effect” between AI Act’s art. 86 and other law including the GDPR’s art. 22).

37. AI Act, *supra* note 2, art. 10(5).

38. *Id.* art. 59(1); see Part II.

39. AI Act, *supra* note 2, arts. 70, 74.

40. See *id.* art. 43(3); see also Part II.

41. See João Pedro Quintais, *Generative AI, Copyright and the AI Act*, 56 *COMPUT. L. & SEC. REV.* 106107 (2025).

C. MEMBER STATE INSTITUTIONS

The Act also relies on several institutions belonging not to the EU, but to member states. Member states are instructed to “establish or designate as national competent authorities at least one *notifying authority* and at least one *market surveillance authority*.”⁴² Recall that the NLF product safety regime on which the Act is built relies on a conformity assessment process that in effect certifies compliance with the law before a product can move on the EU market. Sometimes, conformity assessments are conducted by third parties known as conformity assessment bodies, which are usually private organizations. “Notifying authorities” are member-state-level institutions that can set up procedures to certify conformity assessment bodies as “notified bodies.”⁴³ A conformity assessment body can become a notified body by applying to the notifying authority and satisfying its procedure to check for competence.⁴⁴

The second type of institution, the market surveillance authority, is typically a state agency that oversees product safety generally. The AI Act permits member states to designate other authorities than their existing product safety authorities as market surveillance authorities for purposes of the Act. As discussed below, these member-state-level institutions are important to the accountability frameworks for the Act, and central to the NLF as a whole.

The AI Act also relies on member state institutions that relate to other areas of law. The Act specifically relies on authorities charged with the enforcement of fundamental rights. It grants “the power to request and access any documentation created or maintained under this Regulation” to “[n]ational public authorities or bodies which supervise or enforce the respect of obligations under Union law protecting fundamental rights.”⁴⁵ This description is vague in order to account for the variation in such bodies among member states, and as instructed by the AI Act,⁴⁶ most member states have designated which specific bodies meet that criterion.⁴⁷ To the extent the law intersects with data protection, member state data protection authorities will also have a role in enforcement.

42. AI Act, *supra* note 2, art. 70 (emphases added).

43. *Id.* art. 28(1).

44. *Id.* art. 29.

45. *Id.* art. 77(1).

46. *Id.* art. 77(2).

47. Poklaszlo, *Responsible Authorities for the Enforcement of the AI Act on National Level*, GDPRBLOG (Oct. 2, 2024), https://gdpr.blog.hu/2024/10/02/responsible_authorities_for_the_enforcement_of_the_ai_act_on_national_level (collecting the information).

III. THE LAW

The AI Act centrally aims to promote the uptake of AI throughout Europe, by mitigating risks from the use of AI systems. The very first article of the Act states: “[t]he purpose of this Regulation is to improve the functioning of the internal market and promote the uptake of human-centric and trustworthy artificial intelligence (AI), while ensuring a high level of protection of health, safety, fundamental rights”⁴⁸

For readers outside of the EU accustomed to the stereotype of top-down centralized and heavy-handed regulation, the Act’s explicit prioritization of market functioning and of the uptake of AI may be surprising. But the EU was initially a trading bloc, with its early laws concerned with opening up its common market.⁴⁹ Thus, to read the Act solely as an instantiation of Europe’s adoption of the precautionary principle is a mistake. The Act is in large part concerned with clearing the way for the circulation and adoption of AI through the many member states of EU, through a uniform regulatory framework.⁵⁰ Whether that framework works or not is a different discussion.

The AI Act’s regulatory framework is centrally concerned with mitigating risks.⁵¹ The Act defines “risk” as “the combination of the probability of an occurrence of harm and the severity of that harm.”⁵² It attempts to mitigate multiple kinds of risks through a largely unified approach, categorizing AI into different “risk tiers.” However, the types of risks the Act attempts to mitigate are varied: risks to “health, safety, [and] fundamental rights.”⁵³ While this may appear to be a short list, the reference to “fundamental rights” opens up the Act’s coverage considerably, to include every fundamental right referenced in the Charter.

This attempt to use the same framework for diverse types of risks creates problems. A quantitative framework focused on predicting and measuring potential harms can work well for some types of harms (e.g., to physical safety)

48. AI Act, *supra* note 2, art. 1(1).

49. The precursor to the EU was the European Economic Community (EEC), which began as a pact to regulate coal and steel. *See, e.g., European Union: History*, GLOBALEDGE, <https://globaledge.msu.edu/trade-blocs/european-union/history>.

50. *See* AI Act, *supra* note 2, recital 1 (“This Regulation ensures the free movement, cross-border, of AI-based goods and services, thus preventing Member States from imposing restrictions on the development, marketing and use of AI systems, unless explicitly authorised by this Regulation.”).

51. *See generally* Margot E. Kaminski, *Regulating the Risks of AI*, 103 B.U. L. REV. 1347 (2023).

52. AI Act, *supra* note 2, art. 3(2).

53. *Id.* art. 1(1).

but not for others (e.g., to human rights).⁵⁴ And the type of risk mitigation that might be appropriate for preventing physical harms might look quite different from how one ideally would regulate, for example, harms to privacy or free speech. Moreover, the Act's central reliance on a risk-mitigation (rather than rights-protective) framework arguably assumes AI adoption and that in the process of AI adoption some residual harms to fundamental rights are acceptable.

At its core—its regulation of “high risk” AI systems—the AI Act builds on the framework of EU product safety law. However, multiple aspects of the Act were added ad hoc and thus don't fit squarely into the product safety framework. Some were added at later stages of drafting. For example, after ChatGPT was publicly released during the Act's drafting process,⁵⁵ lawmakers came up with a distinct framework for regulating what the Act terms “general-purpose AI.” Other ad hoc elements arose in response to feedback from various constituencies, including data protection regulators.⁵⁶

Framing the substance of the Act in this way—(A) product safety core, (B) “general-purpose AI” sections, and (C) ad hoc elements—can make it easier to navigate. Most of our discussion in this Part II centers on the Act's product safety core. But we cover each of these three aspects in more detail.

A. THE CORE PRODUCT SAFETY REGIME: THE “NEW LEGISLATIVE FRAMEWORK”

The core of the AI Act, its regulation of high-risk AI systems, adopts the EU's framework approach to product safety.⁵⁷ This “New Legislative Framework” (NLF) has been laid out in two regulations (2008, 2019)⁵⁸ and one directive (2008),⁵⁹ and in product-specific laws.⁶⁰ The procedural core of the

54. *But see, e.g.*, Alessandro Mantelero, BEYOND DATA: HUMAN RIGHTS, ETHICAL AND SOCIAL IMPACT ASSESSMENT IN AI (2022) (arguing for transplanting international law's human rights impact assessment regime into data protection and AI law).

55. Claire Boine & David Rolnick, *Why the AI Act Fails to Understand Generative AI*, 26 MINN. J.L. SCI. & TECH. 61 (2025).

56. *See, e.g.*, Kaminski & Malgieri, *supra* note 36.

57. Michael Veale & Frederik Zuiderveen Borgesius, *Demystifying the Draft EU Artificial Intelligence Act*, 22 COMPUT. L. REV. INT'L 97 (2021); Nicolas Petit & Marco Almada, *The EU AI Act: Between the Rock of Product Safety and the Hard Place of Fundamental Rights*, 62 COMMON MKT. L. REV. 85 (2025).

58. Regulation (EC) No. 765/2008, 2008 O.J. (L 218) 30, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32008R0765>; Regulation (EU) 2019/1020, 2019 O.J. (L 169) 1, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:32019R1020>.

59. Decision 768/2008/EC, 2008 O.J. (L 218) 82, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32008D0768>.

60. *New Legislative Framework*, EUR. COMM'N, https://single-market-economy.ec.europa.eu/single-market/goods/new-legislative-framework_en.

AI Act thus adopts the same framework used for regulating products such as medical devices, construction products, and children's toys.⁶¹

The NLF operates roughly as follows: in order to legally be distributed on the EU market, providers of a product must undergo a "conformity assessment" that indicates the product and its safety governance are in compliance with EU law. After undergoing the conformity assessment, providers of the product may label it as "CE" and then circulate the product on the internal market. Labeled products are then subject to post-market supervision by entities known as "market surveillance authorities." We go through all of this specific to the AI Act in more detail below.

First, we discuss the Act's risk tiers in Section III.A.1. Then we briefly discuss its bans in Section III.A.2. We then, in much greater detail, discuss the Act's regulation of high-risk AI systems, which constitutes the bulk of the Act, and its use of the scaffolding of the New Legislative Framework in Section III.A.3.

1. *The Risk Tiers*

The Act structures its core framework around three tiers of risks raised by particular applications of AI: unacceptable risks, high risks, and everything else. It can be helpful to think of the tiers as a traffic light.⁶² The Act prohibits a specific list of AI practices in Chapter II (red light).⁶³ It regulates a potentially evolving list of AI practices in Chapter III (yellow light).⁶⁴ And it permits everything else, subject to suggested self-regulation (green light).⁶⁵ The bulk of the Act focuses on the regulation of high-risk AI systems (yellow light). We do the same here.

We make one high-level observation before proceeding. First note that the Act regulates two primary sets of actors: AI providers and AI deployers.⁶⁶ (It also regulates AI distributors and importers, which makes sense if you think of the law as being concerned with the introduction of a product onto the EU

61. *Id.*

62. Nicolas Petit points out that the traffic light analogy only goes so far, as an AI system can actually fall into multiple categories at once. But we are trying to produce a simplified guide, so let's stick with the traffic light, for now.

63. AI Act, *supra* note 2, art. 5.

64. *Id.* ch. 3 (High-Risk AI Systems); *see also id.* Annex III.

65. *See, e.g.,* AI Act, *supra* note 2, art. 95(1) (suggesting voluntary Codes of Conduct) ("The AI Office and the Member States shall encourage and facilitate the drawing up of codes of conduct, including related governance mechanisms, intended to foster the voluntary application to AI systems, other than high-risk AI systems, of some or all of the requirements set out in Chapter III, Section 2 taking into account the available technical solutions and industry best practices allowing for the application of such requirements.").

66. *Id.* art. 2(1)(c).

market.⁶⁷) The law's primary focus is on AI providers—that is, what many refer to as AI developers.⁶⁸ It also, however, regulates AI deployers: the users of AI systems.⁶⁹

The core structure of the Act, then, regulates a particular mental model of AI: AI that is developed by one actor and used by another, *for a particular purpose*. The Act's primary obligations attempt to trace accountability as it passes between these two categories of actors.⁷⁰ Which risk tier applies to a particular "AI practice" depends on the specific purpose of the particular AI system. It should come as little surprise, then, that this purpose-centric framing breaks down when it comes to general-purpose AI systems.⁷¹

2. *The Bans*

The AI Act's regulation of the riskiest tier employs a simple but serious legal mechanism: the ban.

Although a short Chapter, the AI Act's regulation of Prohibited AI Practices through bans is in itself rather remarkable. Some uses of AI systems have been banned or paused in the United States, on the municipal level⁷² or the state level.⁷³ There have also been discussions of banning certain uses of AI internationally, in war.⁷⁴ But generally speaking, as of the enactment of the

67. In the interest of relative brevity, we do not cover these obligations in any detail here. Suffice it to say that they consist largely of making sure the AI provider is in compliance with the law, including by checking for conformity markings. *Id.* art. 2(1)(d); *see also id.* arts. 23 (importers)—24 (distributors).

68. This also includes the entity having an AI system developed. *Id.* art. 3(3) ("natural or legal person, public authority, agency or other body that develops an AI system or a general-purpose AI model or that has an AI system or a general-purpose AI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge").

69. *Id.* art. 3(4) ("a natural or legal person, public authority, agency or other body using an AI system under its authority except where the AI system is used in the course of a personal non-professional activity").

70. *See id.* art. 13 (discussing transparency to deployers); *see also id.* art. 14 (enabling human oversight); *see also* Rebecca Crootof, Margot E. Kaminski & W. Nicholson Price II, *Humans in the Loop*, 76 VAND. L. REV. 429, 503 (2023) (discussing the draft AI Act's approach to human oversight). The Act also creates some specific requirements for importers and distributors, but those are designed primarily to prevent loopholes—for example, creating a system abroad without complying. AI Act, *supra* note 2, arts. 23–24.

71. Boine & Rolnick, *supra* note 55.

72. *See* AP, *Seattle Mayor Ends Police Drone Efforts*, USA TODAY, <https://www.usatoday.com/story/news/nation/2013/02/07/seattle-police-drone-efforts/1900785/> (last updated Feb. 7, 2013).

73. S.B. 25-143, Gen. Assemb., Reg. Sess. (Colo. 2025).

74. STOP KILLER ROBOTS, <https://www.stopkillerrobots.org/>; *see* Rebecca Crootof, *The Killer Robots Are Here: Legal and Policy Implications*, 36 CARDOZO L. REV. 1837 (2015) (discussing bans versus regulation).

AI Act, few AI systems have been meaningfully regulated, let alone banned in the United States.

As a matter of regulatory design, the AI Act's use of bans sets a relatively clear set of rules for certain specific AI practices. If regulated companies truly want clarity, a ban on certain practices offers it. A ban is not a standard.⁷⁵ It does not allow (much) room for interpretation in application.⁷⁶ Nor does it delegate decision-making to other bodies, like agencies or courts. The decision of what to ban is made by public lawmakers, not some set of private actors operating through technical standards-setting institutions.⁷⁷ And although much of what the AI Act does is lighter-touch regulation, this component most certainly is not. Recall that the stated purpose of the AI Act is to facilitate the uptake of AI systems on the EU market. The Act's bans of certain uses of AI systems in Chapter II act as an ostensible backstop to that overarching goal of uptake through regulation, offering at least symbolic outer limits to what Europe will allow.⁷⁸

The AI Act bans a list of "Prohibited AI Practices" in the text of the Act itself, rather than an annex.⁷⁹ This makes the list stickier and harder to amend than, say, the Act's list of "high-risk" practices.⁸⁰ It ostensibly reflects a European consensus on which specific uses of AI de facto violate fundamental rights or pose too high a threat to health or safety.⁸¹

75. Louis Kaplow, *Rules v. Standards: An Economic Analysis*, 42 DUKE L.J. 557 (1992). But see Pierre Schlag, RULES AND STANDARDS, 33 UCLA L. REV. 379 (1985).

76. There are exceptions, *see, e.g.*, the exception for uses of emotional inference AI systems in the workplace and educational settings for medical or safety reasons. AI Act, *supra* note 2, art. 5(1)(f).

77. See *id.* arts. 112(1), 112(11) for how amendments to the list of banned practices might happen.

78. See *generally id.* ch. 2.

79. *Id.* art. 5.

80. The default is that the AI Act may be amended through the ordinary EU legislative process, which involves all three of the Commission, Parliament, and Council of the European Union. This takes longer than the streamlined process envisioned, for example, for updating Annex III of the AI Act, discussed below. See AI Act, *supra* note 2, art. 112(1) ("The Commission shall assess the need for amendment of the list set out in Annex III and of the list of prohibited AI practices laid down in Article 5, once a year following the entry into force of this Regulation, and until the end of the period of the delegation of power laid down in Article 97. The Commission shall submit the findings of that assessment to the European Parliament and the Council."). Art. 112(11) envisions the AI Office in the Commission as taking a lead role in coming up with "an objective and participative methodology for the evaluation of risk levels based on the criteria outlined in the relevant Articles and the inclusion of new systems in: . . . (b) the list of prohibited practices set out in Article 5" *Id.*

81. See Ljupcho Grozdanovski & Jérôme De Cooman, *Forget the Facts, Aim for the Rights! On the Obsolescence of Empirical Knowledge in Defining the Risk/Rights-Based Approach to AI Regulation in the European Union*, 49 RUTGERS COMPUT. & TECH. L.J. 207 (2023) (for criticism of the arbitrariness of the list).

In reality, however, the list appears to be a response to a number of salient news stories about the misuse of algorithms or AI systems. Banned practices include: the use of AI for social scoring (the “let’s-not-be China ban”);⁸² the use of AI for predicting the risk of committing a criminal offense (the “Minority Report ban”);⁸³ the use of AI to infer emotions at work or in an educational setting (the “Emotional Phrenology ban”⁸⁴),⁸⁵ and the scraping of public images to contribute to the creation of a facial recognition system (the “Clearview AI ban”)⁸⁶ The full list is outlined in Chapter II, Article 5.⁸⁷

One prohibited practice that merits further discussion is the ban on the use of real-time biometric systems, in public spaces, for law enforcement use (we call it “the biometrics ban” for the rest of this subpart).⁸⁸ That is for two reasons. First, the biometrics ban establishes a ban in close proximity to uses permitted but regulated by other portions of the Act. The Act bans only a very specific subset of uses of biometric systems: real-time (as opposed to after-the-fact), in public spaces (as opposed to in private spaces), and for law enforcement use (as opposed to private sector use).⁸⁹ Otherwise, biometric

82. AI Act, *supra* note 2, art. 5(1)(c).

83. *Id.* art. 5(1)(d).

84. Luke Stark & Jevan Hutson, *Physiognomic Artificial Intelligence*, 32 FORDHAM INTELL. PROP., MEDIA & ENT. L.J. 922 (2021) (calling it phrenology/physiognomy).

85. AI Act, *supra* note 2, art. 5(1)(f) (“the placing on the market, the putting into service for this specific purpose, or the use of AI systems to infer emotions of a natural person in the areas of workplace and education institutions, except where the use of the AI system is intended to be put in place or into the market for medical or safety reasons”).

86. *Id.* art. 5(1)(e).

87. In addition to the named exceptions, *supra* note 76, it includes: the use of subliminal techniques, *id.* art. 5(1)(a), materially distorting persons’ behavior through exploiting known vulnerabilities, *id.* art. 5(1)(b), and the making of certain inferences (race, political opinion) through biometric data, *id.* art. 5(1)(g).

88. *Id.* art. 5(1)(h).

89. See Veale & Zuiderveen Borgesius, *supra* note 57 (“only ‘real-time’ systems that capture, compare, and identify ‘instantaneously, near-instantaneously or in any event without a significant delay’ are prohibited. This excludes ‘post’ systems which, for example, biometrically analyse footage after an event, for example to identify individuals at protests after-the-fact, and systems that categorise individuals biometrically. As online spaces are also out-of-scope, live biometric identification on e.g., video streams is also excluded.”); see also *Cop Out: Security Exemptions in the Artificial Intelligence Act*, STATEWATCH, <https://www.statewatch.org/automating-authority-artificial-intelligence-in-european-police-and-border-regimes/2-cop-out-security-exemptions-in-the-artificial-intelligence-act/> (referring to “(Un)prohibited practices” and observing that “despite supposed bans on practices such as profiling, biometric categorization, and mass biometric surveillance, law enforcement and migration authorities enjoy numerous exemptions that may enable widespread deployment of these techniques”); Francesca Palmiotto, *The AI Act Roller Coaster: The Evolution of Fundamental Rights Protection in the Legislative Process and the Future of the Regulation*, 16 EURO. J. RISK REG. 770, 780, 789 (2025) (“The list of prohibited AI practices in Article 5 has been the most debated

systems are regulated under the Act as high-risk uses of AI.⁹⁰ This establishes a potential cliff effect, whereby regulated AI providers may do their darndest not to fall into the prohibited category of behavior. Some have noted, too, that the Act does not ban but rather condones the installation of the infrastructure for facial recognition.⁹¹

Second, the biometrics ban, which was much-debated during drafting, is the only ban—and indeed, one of the few places in the Act—to envision and structure significant interplay with courts. Unlike several of the other bans, the biometrics ban contains three named exceptions: for a targeted search for victims, for a “specific, substantial and imminent threat to the life or physical safety” or risk of a terrorist attack, and for finding or identifying a person suspected of having committed certain more serious crimes.⁹² It sets up a system of prior judicial oversight (or similar independent administrative oversight, depending on national law) over these exceptions.⁹³ This is notable because it invokes courts directly as the authority on fundamental rights oversight, unlike other aspects of the Act.

3. *High-Risk AI Systems*

Most of the AI Act concerns the regulation of high-risk AI systems—the yellow in our traffic light of risk tiers.⁹⁴ Here, we discuss: which AI systems are considered “high risk,” the substance of the regulation for each of providers

of the AI Act, particularly regarding real-time biometric identification The AI Act embeds double standards for individuals affected by AI systems, with lower protection for individuals suspected or accused of having committed a crime, migrants, asylum seekers and refugees. From a legal and ethical perspective, this is a critical weakness of the Regulation.”).

90. See AI Act, *supra* note 2, Annex III (defining high-risk AI systems to include):

Biometrics, in so far as their use is permitted under relevant Union or national law:

- (a) remote biometric identification systems. This shall not include AI systems intended to be used for biometric verification the sole purpose of which is to confirm that a specific natural person is the person he or she claims to be;
- (b) AI systems intended to be used for biometric categorisation, according to sensitive or protected attributes or characteristics based on the inference of those attributes or characteristics;
- (c) AI systems intended to be used for emotion recognition.

91. Veale & Zuiderveen Borgesius, *supra* note 57, at 102 (“any authorisation of biometrics necessitates installing re-purposable infrastructure. Many already argue the Draft AI Act legitimises rather than prohibits population-scale surveillance”).

92. AI Act, ch. 2, art. 5(1)(h)(i)–(iii); see also *id.* Annex II (listing covered criminal offenses).

93. *Id.* art. 5(3).

94. AI Act, *supra* note 2, ch. III.

(developers) and deployers (users) of high-risk AI systems, how conformity assessment works specific to the Act (including the central role of private technical standards-setting), and the complex accountability systems established by the Act (including registration and post-market monitoring).

a) What Practices Are “High Risk”?

An AI system is considered “high risk” if it is a product or safety component of a product covered by certain named EU laws under the New Legislative Framework⁹⁵ and elsewhere.⁹⁶ Additionally, an AI system is considered “high risk” if it is named in Annex III of the Act.⁹⁷ Annex III currently lists eight categories of high-risk AI systems. These categories include: biometric systems that weren’t banned in the Act and are otherwise permitted under Union and national law;⁹⁸ AI systems used in critical infrastructure;⁹⁹ certain AI systems used in educational and vocational training;¹⁰⁰ certain AI systems used in employment, including in recruitment, termination, and employee monitoring;¹⁰¹ AI systems used for determining certain public and private essential benefits;¹⁰² and three more.¹⁰³

Aiming to future-proof the law, the AI Act makes the list of high-risk AI systems in Annex III easier to update than typical EU law. Article 7 gives the Commission the authority to adopt “delegated acts” to update Annex III, so long as the risk of harm is as high as existing practices on the current list and the new high-risk systems fall under the same existing eight categories.¹⁰⁴

There are two wrinkles of note. First, some of the practices listed in Annex III may in practice be banned, if their use is not permitted by other EU or national law.¹⁰⁵ That is, developers should not assume that an AI practice

95. AI Act, *supra* note 2, art. 6(1). See also *id.* Annex I(A) for the list of relevant EU laws. That list includes: machinery, toys, recreation craft and personal watercraft, elevators, medical devices, and more.

96. See *id.* Annex I(B), which includes railroads, aircraft, and quadricycles, among other things.

97. *Id.* art. 6(2); see also *id.* Annex III.

98. *Id.* Annex III(1).

99. *Id.* Annex III(2).

100. *Id.* Annex III(3).

101. *Id.* Annex III(4).

102. *Id.* Annex III(5).

103. The remaining three are: certain law enforcement uses (Annex III(6)); certain uses in managing asylum, migration, and border control (Annex III(7)); and uses by judges (Annex III(8)(a)) and to influence elections (Annex III(8)(b)).

104. AI Act, *supra* note 2, art. 7.

105. See, e.g., *id.* Annex III(6)–(7) (“in so far as their use is permitted under relevant Union or national law”).

named in Annex III is necessarily legal; rather, it is just not banned by the AI Act specifically.¹⁰⁶

Second, Article 6 contains a significant potential loophole with respect to the Act's high-risk categorization.¹⁰⁷ Providers can, in effect, opt out of having their systems designated as high-risk in some circumstances. Largely, these circumstances involve using AI as a *de minimis* element of an otherwise human decision-making process, even where that decision-making process falls into one of the eight categories in Annex III.¹⁰⁸ This creates incentives to put more humans in the loop for these kinds of decisions, whether authentically or performatively, in order to escape the Act's requirements.¹⁰⁹

A provider may themselves determine under those circumstances that an AI system “does not pose a significant risk of harm to the health, safety or fundamental rights of natural persons, including by not materially influencing the outcome of decision making.”¹¹⁰ There is no *ex ante* scrutiny of this decision. The provider must document this determination and register the

106. *See also* European Data Protection Board-European Data Protection Supervisor, Joint Opinion 5/2021 on the Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) (June 18, 2021), https://www.edpb.europa.eu/our-work-tools/our-documents/edpb-edps-joint-opinion/edpb-edps-joint-opinion-52021-proposal_en [hereinafter EDPB-EDPS Joint Opinion].

107. *See* AI Act, *supra* note 2, art. 6(3). The AI Act again establishes an easier lawmaking process by allowing for amendment through delegated acts “by adding new conditions to those laid down therein, or by modifying them, where there is concrete and reliable evidence of the existence of AI systems that fall under the scope of Annex III, but do not pose a significant risk of harm to the health, safety or fundamental rights of natural persons.” *Id.* art. 6(6).

108. *Id.* art. 6(3) states that the exception “shall apply where *any* of the following conditions is fulfilled:

- (a) the AI system is intended to perform a narrow procedural task;
- (b) the AI system is intended to improve the result of a previously completed human activity;
- (c) the AI system is intended to detect decision-making patterns or deviations from prior decision-making patterns and is not meant to replace or influence the previously completed human assessment, without proper human review; or
- (d) the AI system is intended to perform a preparatory task to an assessment relevant for the purposes of the use cases listed in Annex III.”

(emphasis added).

109. *See* Crotoft et al., *supra* note 70, at 450 (“Some laws encourage regulatory arbitrage when retaining a human in the loop allows a system’s developers or users to avoid more onerous regulation.”).

110. AI Act, *supra* note 2, art. 6(3). However, a provider cannot self-designate as not-high-risk an AI system that “performs profiling of natural persons.”

system before putting it on the market, so that the determination can be examined after the fact if harms should occur or contrary evidence should arise.¹¹¹

b) What Are the Substantive Requirements for Providers
(Developers) of High-Risk AI Systems?

The bulk of the AI Act aims to regulate providers of high-risk AI systems. Most of these requirements are outlined in Chapter III, Section 2 (helpfully titled “Requirements for high-risk AI systems”).¹¹² Requirements include, among other things: establishing a risk mitigation system;¹¹³ establishing data governance;¹¹⁴ establishing technical documentation¹¹⁵ and automated logging;¹¹⁶ and establishing certain accuracy, robustness, and cybersecurity measures.¹¹⁷

Some of these requirements are more substantive. Others primarily establish procedures and documentation towards enabling later external accountability. Some entwine procedure with substance. And several are aimed at maintaining accountability during a handoff from system developer to system deployer.¹¹⁸

This section, like much of this Article, is not intended to provide legal advice nor to be exhaustive. We provide an overview of the Act’s requirements that apply to providers of high-risk AI systems, with several more in-depth examples of each kind of requirement.

First, the AI Act, for all its procedural and accountability scaffolding, does contain substantive requirements. There are several examples in Article 15, on “Accuracy, robustness and cybersecurity.”¹¹⁹ The Act, for example, requires high-risk AI systems to “achieve an appropriate level of accuracy,”¹²⁰ and refers to developing substantive “benchmarks and measurement methodologies” as performance metrics.¹²¹ The Act requires that developers disclose said levels of accuracy and metrics in instructions conveyed to deployers.¹²²

111. *Id.* art. 6(4). National competent authorities may later ask for this documentation.

112. *Id.* ch. III, sec. 2.

113. *Id.* art. 9.

114. *Id.* art. 10.

115. *Id.* art. 11.

116. *Id.* arts. 12, 19 (obligation to retain logs).

117. *Id.* art. 15.

118. *See, e.g., id.* art. 14 (on human oversight), *id.* art. 13 (on transparency to deployers).

119. *Id.* art. 15.

120. *Id.* art. 15(1).

121. *Id.* art. 15(2).

122. *Id.* art. 15(3).

Other examples of substantive requirements can be found in Article 10 on data and data governance.¹²³ The Act requires that “[t]raining, validation and testing data sets shall be relevant, sufficiently representative, and to the best extent possible, free of errors and complete in view of the intended purpose.”¹²⁴ It additionally requires that such data sets must “have the appropriate statistical properties . . . as regards the persons or groups of persons in relation to whom the high-risk AI system is intended to be used.”¹²⁵ The latter requirement is significant, given high-profile reports of algorithms developed using data from one population but then used on another, without confirming that the two populations had the same statistical properties.¹²⁶

Second, elements of this Section of the Act require documentation and process towards establishing external accountability. For example, Article 11 requires technical documentation.¹²⁷ The elements of such technical documentation are outlined in Annex IV, and smaller businesses may adopt a simplified version.¹²⁸ The explicit purpose of the requirement of technical documentation is to establish external accountability to government actors.¹²⁹

Third, elements of this Section of the Act involve process mixed with substance. For example, Article 9 requires that developers establish a “risk management system” for high-risk AI systems.¹³⁰ At first glance, these requirements look primarily procedural, requiring a set of steps.¹³¹ But the risk

123. *Id.* art. 10.

124. *Id.* art. 10(3).

125. *Id.* arts. 10(3)–(4) (“Data sets shall take into account, to the extent required by the intended purpose, the characteristics or elements that are particular to the specific geographical, contextual, behavioural or functional setting within which the high-risk AI system is intended to be used.”).

126. *See* State v. Loomis, 2016 WI 68, ¶¶ 63–66 (noting that COMPAS, which was trained on data from Broward County, Florida, and applied in Wisconsin, had not yet had a cross-validation study for the Wisconsin population). For more examples, see Angelina Wang, Sayash Kapoor, Solon Barocas & Arvind Narayanan, *Against Predictive Optimization: On the Legitimacy of Decision-Making Algorithms That Optimize Predictive Accuracy*, 1 ACM J. RESPONSIBLE COMPUTING 9:1, 9:12–13 (Mar. 20, 2024) (discussing examples of “distribution shifts”).

127. AI Act, *supra* note 2, art. 11.

128. *Id.* art. 11(1); *see also id.* Annex IV.

129. *Id.* art. 11(1) (“demonstrate that the high-risk AI system complies with the requirements set out in this Section and to provide national competent authorities and notified bodies with the necessary information in a clear and comprehensive form to assess the compliance of the AI system with those requirements.”).

130. *Id.* art. 9(1) (“A risk management system shall be established, implemented, documented and maintained in relation to high-risk AI systems.”).

131. Article 9(2) of the AI Act states that:

(a) the identification and analysis of the known and the reasonably foreseeable risks that the high-risk AI system can pose to health, safety or

management system also has substantive aspects. For example, it requires mitigation of risks down to an “acceptable” level of risk, both for each hazard and for the system overall.¹³² Article 9 also requires that high-risk AI systems be tested.¹³³ Substantively, it requires these systems be fit for their intended purpose.¹³⁴

Fourth, several of the Act’s requirements for developers aim to establish an accountability hand-off from providers (developers) to deployers (users) of high-risk AI systems. Article 13 requires that system developers create and transmit instructions for and transparency to deployers.¹³⁵ Instructions must include at least “the characteristics, capabilities and limitations of performance” of the AI system, including its purpose; its level of accuracy, robustness, and cybersecurity; circumstances that may lead to higher risks; and “its performance regarding specific persons or groups of persons on which the system is intended to be used.”¹³⁶ The Act also requires explanations of the AI systems to deployers (as opposed to affected individuals, which is addressed in Article 86), for example requiring that providers disclose “where applicable, the technical capabilities and characteristics of the high-risk AI system to provide information that is relevant to explain its output” to deployers.¹³⁷

In some places, however, the envisioned accountability hand-off may break down. Take, for example, the provisions on human oversight, which

fundamental rights when the high-risk AI system is used in accordance with its intended purpose;

(b) the estimation and evaluation of the risks that may emerge when the high-risk AI system is used in accordance with its intended purpose, and under conditions of reasonably foreseeable misuse;

(c) the evaluation of other risks possibly arising, based on the analysis of data gathered from the post-market monitoring system referred to in Article 72;

(d) the adoption of appropriate and targeted risk management measures designed to address the risks identified pursuant to point (a).

132. *Id.* art. 9(5).

133. *Id.* art. 9(6) (“High-risk AI systems shall be tested for the purpose of identifying the most appropriate and targeted risk management measures. Testing shall ensure that high-risk AI systems perform consistently for their intended purpose and that they are in compliance with the requirements set out in this Section.”).

134. *Id.* (“Testing shall ensure that high-risk AI systems perform consistently for their intended purpose”).

135. *Id.* arts. 13(1)–(2).

136. *Id.* art. 13(3).

137. *Compare id.* art. 13(3)(b)(vii) (“where applicable, information to enable deployers to interpret the output of the high-risk AI system and use it appropriately”), *with* GDPR, *supra* note 19, art. 15(h) (requiring disclosure to affected individuals of “meaningful information about the logic involved”); *see also* Case C-203/22, *supra* note 36 (interpreting art. 15 of GDPR in light of art. 22).

describe a complex pass-off from providers to deployers. The Act requires that high-risk AI systems be designed by providers for human oversight “commensurate with the risks, level of autonomy and context of use.”¹³⁸ AI providers are responsible for building related enabling features into AI.¹³⁹ Providers are also responsible for telling deployers how to use high-risk AI, including how much human oversight is needed.¹⁴⁰ However, the Act recognizes that human oversight itself will often be implemented by deployers.¹⁴¹ Thus, the envisioned pass-off ideally goes as follows: AI providers design their systems for human oversight (as a form of risk mitigation) and instruct deployers as to how much oversight should be used. Then, deployers must follow the instructions.

The Act however, largely declines to put direct obligations on the deployers.¹⁴² As written, the Act delegates to AI providers to decide what level of human oversight is required, and to govern deployers through instructions for use.¹⁴³ Michael Veale and Frederik Zuiderveen Borgesius describe the set-up as follows: “Somewhat strangely, no obligations for human oversight flow directly from the Act to a user. In relation to human oversight, users must simply follow the instruction manual.”¹⁴⁴ The only direct obligations on deployers related to human oversight are to assign it to “natural persons who have the necessary competence, training and authority, as well as the necessary support,”¹⁴⁵ and to use high-risk AI systems “in accordance with the

138. AI Act, *supra* note 2, art. 14(3).

139. *Id.* art. 14(3)(a).

140. *Id.* art. 13(3)(d) (“the human oversight measures referred to in Article 14, including the technical measures put in place to facilitate the interpretation of the outputs of the high-risk AI systems by the deployers”).

141. *Id.* art. 14(3)(b) (“measures identified by the provider before placing the high-risk AI system on the market or putting it into service and that are appropriate to be implemented by the deployer”).

142. See Veale & Zuiderveen Borgesius, *supra* note 57, at 104 (“Somewhat strangely, no obligations for human oversight flow directly from the Act to a user. In relation to human oversight, users must simply follow the instruction manual.”); see also Crootof et al., *supra* note 70, at 503–04 (“The Act divides regulated entities into providers, who build AI systems, and users, who use them. As a consequence, *nobody is really responsible for the human-machine system as a whole.*”).

143. There is one notable exception: the AI Act requires that a deployer may not take action on the basis of identification by remote biometric identification systems unless at least two humans “with the necessary competence, training and authority” have separately confirmed the identification. AI Act, *supra* note 2, art. 14(5). The Act exempts, however, “high-risk AI systems used for the purposes of law enforcement, migration, border control or asylum, where Union or national law considers the application of this requirement to be disproportionate.” *Id.*

144. Veale & Zuiderveen Borgesius, *supra* note 57, at 104.

145. AI Act, *supra* note 2, art. 26(2).

instructions for use.”¹⁴⁶ Some—but not all—deployers must provide a “description of the implementation of human oversight measures, according to the instructions for use.”¹⁴⁷

c) What Are The Substantive Requirements for Deployers (Users) of High-Risk AI Systems?

This leads us into a much briefer discussion of the Act's requirements for deployers (users) of high-risk AI systems. Largely, these are laid out in Article 26.¹⁴⁸ But deployers should be sure to read the Act as a whole, as several requirements for deployers are laid out elsewhere.¹⁴⁹ We outline a few core requirements here.

As noted above, deployers are legally obligated to follow the instructions for use that they receive from providers, including any instructions regarding human oversight.¹⁵⁰ They are required to ensure that input data is “relevant and sufficiently representative in view of the intended purpose of the high-risk AI system.”¹⁵¹ They are required to keep logs automatically generated by AI systems, typically for at least six months.¹⁵² And, centrally, deployers are required to “monitor the operation of the high-risk AI system” and notify the provider, distributor, and relevant government authority when things go wrong and the system presents a risk even when instructions are followed, or when there has been a “serious incident.”¹⁵³ They are required in those cases to stop using the AI system.¹⁵⁴

Deployers are also assigned several notification requirements. They must notify workers and their representatives before deploying high-risk AI systems in the workplace.¹⁵⁵ And they must notify affected individuals if they have been subject to a decision made by a high-risk AI system.¹⁵⁶ Deployers are also subject to several requirements outlined elsewhere in the Act, which we discuss further below: to provide individual explanations of AI decisions,¹⁵⁷ and to

146. *Id.* art. 26(1).

147. Only if they are required to conduct a Fundamental Rights Impact Assessment. *See id.* art. 27(1)(e).

148. *Id.* art. 26.

149. *See, e.g.*, our discussion of “ad hoc” elements below, and discussion of biometric verification in note 143.

150. AI Act, *supra* note 2, art. 26(1).

151. *Id.* art. 26(4).

152. *Id.* art. 26(6).

153. *Id.* art. 26(5).

154. *Id.*

155. *Id.* art. 26(7).

156. *Id.* art. 26(11).

157. *Id.* art. 86

conduct Fundamental Rights Impact Assessments in some contexts, before deploying a system.¹⁵⁸

d) Conformity Assessments

As mentioned, the key feature of the NLF that the AI Act borrows is the “conformity assessment” and the accountability framework that accompanies it.¹⁵⁹ A conformity assessment is a document that states whether the “essential requirements” of the law in question are met, as well as different procedural requirements that the law may implement.¹⁶⁰ In the AI Act, a conformity assessment requires that three categories of requirements be satisfied.¹⁶¹

First is the quality management system. Article 17 lays out thirteen requirements for the quality management system, among them “a strategy for regulatory compliance”; “techniques, procedures and systematic actions” for design, design control, design verification, development, quality control, and quality assurance; “examination, test and validation procedures to be carried out before, during and after the development of the high-risk AI system”; “technical specifications, including standards, to be applied”; comprehensive “systems and procedures for data management”; “the risk management system referred to in Article 9”; “the setting-up, implementation and maintenance of a post-market monitoring system”; “procedures related to the reporting of a serious incident”; and “an accountability framework setting out the responsibilities of the management and other staff with regard to all the aspects listed.”¹⁶²

The second requirement that must be satisfied is that technical documentation must demonstrate compliance with the essential requirements in Chapter III, Section 2 of the law.¹⁶³ Chapter III, Section 2 of the law corresponds to Articles 8–15, so this entails the risk management framework, data governance system, documentation, logging, instructions for use, human oversight, and accuracy and robustness checks described above.¹⁶⁴

The third requirement that must be satisfied is verification of a post-market monitoring system described in Article 72. This is a system of monitoring for data that allows a provider to evaluate continuous compliance

158. *Id.* art. 27.

159. *Id.* art. 43.

160. *Id.* art. 3(20).

161. *Id.* Annexes VI–VII.

162. *Id.* art. 17(1).

163. *Id.* art. 3(20).

164. *Id.* arts. 9–15.

with the regulation throughout the lifetime of the AI system.¹⁶⁵ The Commission is charged with establishing a template for this by early 2026.¹⁶⁶

There are two sets of procedures for undergoing a conformity assessment. One is a self-assessment laid out in Annex VI, and the other is a third-party assessment described in Annex VII.¹⁶⁷ Article 43 dictates which procedures each type of high-risk AI system is subject to. It divides high-risk AI systems into three categories with respect to whether they (a) can opt for either self-certification or third-party assessment, (b) may use the self-certification procedure, or (c) must use third-party assessment.

The bottom line is that (1) providers of biometric systems that follow published standards can self-certify, but if they do not, they must get third-party certification;¹⁶⁸ (2) providers of AI for critical infrastructure and other fundamental-rights-impacting AI can self-certify;¹⁶⁹ and (3) providers of safety-impacting AI—which is safety-impacting only because it is one of the types of products *otherwise* designated as safety-impacting under the NLF—must undergo the third-party assessment procedure dictated by the other product safety law the product is subject to.¹⁷⁰ The fact that rights-impacting AI is governed by self-certification of compliance has led to some of the serious critiques by European scholars and activists that this law does not take fundamental rights seriously enough.¹⁷¹

165. *Id.* art. 72(2).

166. *Id.* art. 72(3).

167. *Id.* art. 43.

168. *Id.* art. 43(1). Under Article 43(1), providers of “high-risk AI systems listed in point 1 of Annex III,”—otherwise known as the biometric systems that are not banned outright—can opt for self-certification or third-party assessment if they’ve followed harmonized standards under Article 40 or the common procedures under Article 41 (more on both below). However, if such standards do not exist or were not followed, the provider must use the Annex VII (third-party assessment) procedure.

169. *Id.* art. 43(2). Article 43(2) says that providers of “high-risk AI systems referred to in points 2 to 8 of Annex III”—otherwise known as the non-biometric rights-impacting AI (points 3–8) plus critical infrastructure (point 2)—may use the Annex VI (self-certification) procedure.

170. *Id.* art. 43(3). Article 43(3) says that providers of “high-risk AI systems covered by the Union harmonisation legislation listed in Section A of Annex I”—otherwise known as safety-impacting AI systems already covered by a different NLF law—should follow the conformity assessment procedures of those other laws and just add the requirements of the AI Act to those conformity assessments.

171. See Daniel Leufer, Fanny Hidvegi & Alessia Zornetta, *The Pitfalls of the European Union's Risk-Based Approach to Digital Rulemaking*, 71 UCLA L. REV. DISCOURSE 156 (2024). We would add that while a cynic—or an American—might not be entirely surprised by fundamental rights being pushed aside, we find the idea that self-certification also applies to critical infrastructure in some ways even more surprising, leading us to wonder whether there is a more fundamental respect for law in play in Europe, such that we should not assume law

While the Act contains these requirements for conformity assessments, in practice many of the details will end up being developed by private standards-setting organizations. Article 40 directs the Commission to issue standardization requests to the ESOs to cover all substantive requirements of the law. Article 40(1) then creates a kind of safe harbor: it states that “[h]igh-risk AI systems or general-purpose AI models which are in conformity with harmonised standards or parts thereof . . . shall be presumed to be in conformity” with the requirements of the law.¹⁷²

Thus, providers can avoid the sometimes vague and difficult questions about when and whether their process lines up with the requirements of the law by adhering strictly to any published standards. The Commission has requested that a joint technical committee of CEN and CENELEC take up the task, which is not yet completed.¹⁷³ Assuming they finish, and the standards meet the criteria in the law,¹⁷⁴ then those standards are poised to become the primary way that companies comply with the AI Act.¹⁷⁵

e) Accountability and Enforcement

The AI Act’s envisioned accountability framework for high-risk systems is based on (a) establishing internal corporate governance, (b) creating documentation that overseeing institutions can examine ex post, (c) setting up systems that will alert overseeing institutions if things go wrong, and (d) enabling market surveillance authorities to obtain access to information and affording them the ability to take remedial or punitive action.

The first component, corporate governance, is accomplished via the Chapter III, Section 2 articles described in the last two parts—for example, creation of risk management systems (Art. 9); data governance systems (Art. 10); human oversight (Art. 14); and checks on accuracy, robustness, and security (Art. 15).¹⁷⁶ The conformity assessment, whether self-certified or

must be written with Holmes’ bad man in mind. On the other hand, the EDPB and EDPS called for third-party oversight and did not succeed in getting it into the law. *See* EDPB-EDPS Joint Opinion, *supra* note 106.

172. AI Act, *supra* note 2, art. 40(1).

173. *Artificial Intelligence*, *supra* note 24; *Commission Implementing Decision on a Standardisation Request to the European Committee for Standardisation and the European Committee for Electrotechnical Standardisation in support of Union Policy on Artificial Intelligence*, at 5, C(2023) 3215 final (May 22, 2023), [https://ec.europa.eu/transparency/documents-register/detail?ref=C\(2023\)3215&lang=en](https://ec.europa.eu/transparency/documents-register/detail?ref=C(2023)3215&lang=en).

174. Shortcomings in either the standards or common specifications are a reason that market surveillance authorities may reject the presumption. AI Act, *supra* note 2, art. 79(6).

175. If the standards are unfinished or inadequate, under Article 41, the Commission may adopt “common specifications” that functionally take the place of technical standards. *See, e.g., id.* art. 43(1) (referring to both Articles 40 and 41 as alternatives to each other).

176. *See* Sections II.D(2)–(3).

certified by a notified body, requires an assessment that these systems have been adequately set up. Similarly, the conformity assessment requires a statement of compliance with Article 17, which requires a “quality management system in place that ensures compliance.”¹⁷⁷

The Act’s second component of accountability and enforcement is documentation. Article 11 requires technical documentation to be kept up to date throughout the life of the product.¹⁷⁸ Article 17 requires that the quality management system “be documented in a systematic and orderly manner in the form of written policies, procedures and instructions.”¹⁷⁹ Both are necessary components of the corporate governance framework. But the real point of the documentation is enabling later oversight by government bodies. Article 18 requires the technical and quality management system documentation be kept for ten years “at the disposal of the national competent authorities,”¹⁸⁰ and Article 21 requires that providers give a requesting authority “all the information and documentation necessary to demonstrate the conformity of the high-risk AI system”¹⁸¹ Article 74 grants “full access by providers to the documentation as well as the training, validation and testing data sets used for the development of high-risk AI systems.”¹⁸² Similarly, Article 77 grants national authorities overseeing fundamental rights access to “any documentation created or maintained under this Regulation in accessible language and format when access to that documentation is necessary for effectively fulfilling their mandates within the limits of their jurisdiction.”¹⁸³

The third component of the Act’s system of accountability and enforcement is affirmative notice to government actors. Article 20 begins with self-governance, requiring that whenever a provider learns that a system is no longer in conformity, they must take immediate corrective action “to bring that system into conformity, to withdraw it, to disable it, or to recall it, as appropriate.”¹⁸⁴ But importantly, it also states that where a system “present[s] risks to the health or safety, or to fundamental rights, of persons”¹⁸⁵ and the provider becomes aware of the risk, it must investigate and, along with the

177. AI Act, *supra* note 2, art. 17(1)(a).

178. *Id.* art. 11(1).

179. *Id.* art. 17(1).

180. *Id.* art. 18(1).

181. *Id.* art. 21(1).

182. *Id.* art. 74(12). A similar provision in Article 91 grants the Commission access to documentation of general-purpose models, because the Commission is designated as the market surveillance authority for general-purpose AI. *See id.* arts. 91(1), 75(1); *see also infra* Section III.C.

183. AI Act, *supra* note 2, art. 77(1).

184. *Id.* art. 20(1).

185. *Id.* art. 79(1) (cited by art. 20(2)).

deployer of the system, inform the market surveillance authorities and notified body (where applicable) of the risk.¹⁸⁶ Combined with the requirements for automatic logging, this provision is designed to alert the government when anything goes wrong.

Finally, the fourth component of the Act's accountability framework is the ability of relevant government bodies to conduct investigations and take remedial and punitive action where necessary. This constitutes a framework of responsive regulation that is common among EU regulators.¹⁸⁷

Article 79 instructs market surveillance authorities to evaluate systems presenting a risk to health, safety, or to fundamental rights once they are made aware of such a risk, and, if appropriate, to cooperate with national authorities that protect fundamental rights.¹⁸⁸ If the authority finds noncompliance, it must require the "relevant operator to take all appropriate corrective actions to bring the AI system into compliance, to withdraw the AI system from the market, or to recall it."¹⁸⁹ If the operator does not take the corrective action, the authority must escalate the matter, prohibiting the product or recalling it, as necessary.¹⁹⁰ Rules about "penalties and other enforcement measures, which may also include warnings and non-monetary measures," are delegated to member states.¹⁹¹

Article 99 provides for varied maximum fines depending on the violations. Violations of the bans in Article 5 are the highest (the higher of €35 million or 7 percent of annual revenue), with noncompliance with certain "provisions related to operators or notified bodies" (€15 million or 3 percent) in the middle, and probably the most common issue—supplying "incorrect, incomplete or misleading information to notified bodies or national competent authorities in reply to a request" (€7.5 million or 1 percent)—the lowest fine, but still substantial.¹⁹²

Overall, this is a relatively light-touch approach to accountability and oversight, that gives primary control to the system providers and relies heavily on good-faith substantive compliance. The clearest example of this is the

186. *Id.* art. 20(2).

187. See William McGeveran, *Friendship the Privacy Regulators*, 58 ARIZ. L. REV. 959, 983–85 (2016); Margot E. Kaminski, *Binary Governance: Lessons from the GDPR's Approach to Algorithmic Accountability*, 92 S. CAL. L. REV. 1529, 1596–97 (2019); see also Dennis D. Hirsch, *Going Dutch? Collaborative Dutch Privacy Regulation and the Lessons It Holds for U.S. Privacy Law*, 2013 MICH. ST. L. REV. 83, 124 (2013) (on the generally collaborative approach of the pre-GDPR Dutch data protection regime).

188. AI Act, arts. 79(1)–(2).

189. *Id.* art. 79(2).

190. *Id.* art. 79(5).

191. *Id.* art. 99(1).

192. *Id.* arts. 99(3)–(5).

ability of a provider to designate a system as *not* high-risk despite its being a system on the list of high-risk types of systems.¹⁹³ The only requirements for opting the system out of the substantive regulations of this law almost entirely are to document the decision, to make that documentation available to the relevant national authorities upon request, and to register it in an EU database of high-risk systems.¹⁹⁴

Other models of oversight are also notably absent. There is no private right of action for injured individuals; instead, the Act provides for a right to lodge a complaint with a market surveillance authority, which is then folded into the authority's standard process.¹⁹⁵ Transparency to the public, too, is limited. While there is extensive documentation, it is intended for government actors, not the public. The Act does provide for a public database of existing high-risk AI systems under Annex III (the rights-impacting plus critical infrastructure systems),¹⁹⁶ but nothing else.

The Act's accountability framework overall thus relies on the hope that AI providers know best how to judge their systems' risks, and that the government should largely monitor, getting involved only if and after things go wrong.¹⁹⁷ Again, this is not really a precautionary regulatory framework.

B. GENERAL-PURPOSE AI MODELS

The Act's approach to what it terms "general-purpose AI" models is distinct from the main portion of the law concerned with "high-risk AI" systems and works differently. It, too, adopts elements of risk regulation. But unlike the Act's approach to predictive AI, it is not built as squarely on the NLF.

The reason for this is in some sense a practical one: The original AI Act, including the reliance on the NLF, was drafted in 2021, before ChatGPT was released to the public and became a household name. When concerns over generative AI exploded in 2022, the law's drafters—who were writing a law

193. See *supra* notes 109–111 and accompanying text.

194. AI Act, *supra* note 2, art. 6(4). Recent proposals to revise the AI Act aim to get rid of this registration condition. See Maria Niestadt (Eur. Parl. Rsch. Serv.), *Briefing: Digital Omnibus on AI*, 3, PE 782.651 (Feb. 2026), [https://www.europarl.europa.eu/RegData/etudes/BRIE/2026/782651/EPRS_BRI\(2026\)782651_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2026/782651/EPRS_BRI(2026)782651_EN.pdf).

195. *Id.* art. 85.

196. *Id.* art. 71.

197. See Veale & Zuiderveen Borgesius, *supra* note 57, at 102 ("The philosophy of the NLF is that '[t]he manufacturer, having detailed knowledge of the design and production process, is best placed to carry out the complete conformity assessment procedure. Conformity assessment should therefore remain the obligation of the manufacturer alone.' This distinguishes NLF regimes (including the Draft AI Act) from pharmaceutical regulation, where a public authority (e.g., the European Medicines Agency) carries out an assessment themselves before granting pre-marketing approval.").

purporting to cover all of AI—were forced to go back and add in some treatment of generative, or what they call “general-purpose,” AI.¹⁹⁸

Rather than rewrite the bill from the ground up, the drafters inserted new provisions that mimicked some of the concepts in the rest of the bill but used different enforcement mechanisms and institutions. Probably the biggest overall difference is that the Act treats these general-purpose AI models as inherently transnational, eschewing reliance on national market surveillance authorities and centralizing oversight at the EU level.

The interaction of these provisions with the rest of the Act is complicated. A general-purpose AI model might end up integrated into a general-purpose AI system, which are not assigned to any of the three risk tiers that make up the rest of the Act.¹⁹⁹ Or, it might end up integrated into an AI system that does fall into one of the Act’s risk tiers.²⁰⁰ In any event, the Act’s general-purpose AI obligations outlined below fall on the providers of general-purpose AI models regardless of whether the system into which the model is integrated fits neatly into the rest of the regulation.

The Act contrasts generative AI models with other types of AI models based primarily on the generality of application, defining a “general-purpose AI model” as a model “that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications.”²⁰¹ The definition “includ[es] where such an AI model is trained with a large amount of data using self-supervision at scale,” but does not limit the definition to such technical parameters. It

198. Council of the EU Press Release 1008/22, Artificial Intelligence Act: Council Calls for Promoting Safe AI that Respects Fundamental Rights (Dec. 6, 2022) (noting new provisions about general-purpose AI being added on December 6, 2022).

199. *See, e.g.*, AI Act, *supra* note 2, recital 100 (“When a general-purpose AI model is integrated into or forms part of an AI system, this system should be considered to be general-purpose AI system when, due to this integration, this system has the capability to serve a variety of purposes. A general-purpose AI system can be used directly, or it may be integrated into other AI systems.”).

200. *See id.* recital 97 (“The notion of general-purpose AI models should be clearly defined and set apart from the notion of AI systems to enable legal certainty Although AI models are essential components of AI systems, they do not constitute AI systems on their own. AI models require the addition of further components, such as for example a user interface, to become AI systems. AI models are typically integrated into and form part of AI systems. This Regulation provides specific rules for general-purpose AI models and for general-purpose AI models that pose systemic risks, which should apply also when these models are integrated or form part of an AI system.”).

201. *Id.* art. 3(63). A “general-purpose AI system,” means “an AI system which is based on a general-purpose AI model and which has the capability to serve a variety of purposes.” *Id.* art. 3(66).

explicitly excludes “AI models that are used for research, development or prototyping activities before they are placed on the market.”²⁰²

The Act separates general-purpose AI models into two distinct categories: those with and without “systemic risk.” A general-purpose AI model has systemic risk if it either “has high impact capabilities evaluated on the basis of appropriate technical tools and methodologies, including indicators and benchmarks” or the Commission sees the model as equivalent to that.²⁰³ “High impact capabilities,” in turn, is a phrase defined in the law as “capabilities that match or exceed the capabilities recorded in the most advanced general-purpose AI models.”²⁰⁴ When the Commission decides whether a general-purpose model poses equivalent systemic risks, it must “take into account” criteria laid out in Annex XIII, which are largely technical criteria, including the size of the model, the amount of computation used to train the model, the different modalities of the model (text-to-text, image-to-text, etc.), and concerns about the model’s reach—namely, the number of registered end-users.²⁰⁵ The Act separately creates a presumption of systemic risk when over 10^{25} floating point operations (FLOPs) are used for training.²⁰⁶ Annex XIII also creates a presumption where the model is available to ten thousand registered business users.²⁰⁷ The Commission has the power to amend these thresholds over time.²⁰⁸

Providers of general-purpose AI models have certain obligations, with additional obligations if the AI model is deemed to have systemic risk. Providers must keep up-to-date technical documentation that can be reviewed by the AI Office and national authorities. They must also make information and documentation available to downstream providers of AI systems who intend to integrate the general-purpose AI model into their system.²⁰⁹ These requirements are similar to the documentation requirements for high-risk

202. *Id.* art. 3(63).

203. *Id.* art. 51(1).

204. *Id.* art. 3(64).

205. *Id.* art. 51(1); *see also id.* Annex XIII. Article 52 contains more detail about the procedure of such a designation and appeals by the providers.

206. *Id.* art. 51(2). Notably, this FLOP-based threshold is similar to one in the Biden Executive Order on AI, except the number there is 10^{26} FLOPs, an order of magnitude higher. This means many more models meet the threshold in Europe than did in the US under Biden. Epoch.AI has studied documentation of more than 300 models, and has, as of May 12, 2025, found one model that would exceed the 10^{26} threshold, while about 20 exceed the 10^{25} line. EPOCH AI, DATA ON AI: AI MODELS (Aug. 12, 2025), <https://epoch.ai/data/ai-models>.

207. AI Act, *supra* note 2, Annex XIII.

208. *Id.* art. 51(3).

209. *Id.* art. 53(1)(b).

systems that must be made available for oversight, as well as the instructions to deployers that are needed to properly use the AI system.²¹⁰

In addition, the Act requires providers of general-purpose AI models to do three things related to copyright law: (1) establish a copyright policy; (2) operationalize rightsholders' opt-outs from training datasets; and (3) draw up and make publicly available a sufficiently detailed summary about content used for training their model.²¹¹ The Commission recently released a template for describing training data.²¹² The Act exempts providers of open-source models from the requirements for general-purpose AI systems, but not from the copyright requirements.

Providers of general-purpose AI models with systemic risk must additionally perform model evaluation, conduct and document adversarial training, assess and mitigate possible systemic risks, report and document information about "serious incidents," and attend to cybersecurity protection.²¹³ The Act does not exempt open-source models with systemic risk.²¹⁴

The Act's oversight framework for general-purpose AI models is distinct from that used for high-risk AI systems. (However, recall that a general-purpose model can end up used in a high-risk AI system, which we discuss below.²¹⁵) For one, the conformity assessment framework—the central framework for the law overall—just doesn't apply to general-purpose AI models.²¹⁶ Instead the Act grants "exclusive powers to supervise and enforce" the relevant provisions to the AI Office of the Commission.²¹⁷ To support this,

210. See *supra* Section III.A.

211. AI Act, *supra* note 2, art. 53(1)(c) (citing Article 4(3) of Directive (EU) 2019/790, which is about the opt-out right).

212. See EUR. COMM'N, TEMPLATE FOR THE PUBLIC SUMMARY OF TRAINING CONTENT FOR GENERAL-PURPOSE AI MODELS (July 24, 2025), <https://ec.europa.eu/newsroom/dae/redirection/document/118578>; *Drawing-Up a General-Purpose AI Code of Practice*, EUR. COMM'N (Aug. 1, 2025), <https://digital-strategy.ec.europa.eu/en/policies/ai-code-practice>; EUROPEAN AI OFFICE, EUROPEAN AI OFFICE WORKING GROUP MEETINGS: CODE OF PRACTICE FOR GENERAL-PURPOSE AI: TEMPLATE FOR SUMMARY OF TRAINING DATA, EUR. COMM'N (Jan. 17, 2025), <https://ec.europa.eu/newsroom/dae/redirection/document/111909>; *Commission Presents Template for General-Purpose AI Model Providers to Summarise the Data Used to Train Their Model*, EUR. COMM'N: PRESS RELEASE (July 24, 2025), <https://digital-strategy.ec.europa.eu/en/news/commission-presents-template-general-purpose-ai-model-providers-summarise-data-used-train-their>.

213. AI Act, *supra* note 2, art. 55.

214. *Id.* art. 53(2).

215. See *also id.* recital 97.

216. Oddly, Articles 40 and 41 include statements that if general-purpose AI systems follow the standards or common specifications, respectively, they are presumed to be in conformity, despite the law nowhere *requiring* them to be in conformity. *Id.* arts. 40(1), 41(1).

217. *Id.* art. 88(1).

the Commission is granted the powers to request documentation and information, to evaluate models, to request remedial measures, and to make them binding after a “structured dialogue” with the provider.²¹⁸

The Act also considers situations governing the overlap of general-purpose models and high-risk AI systems. Where any AI system is made by a provider of a general-purpose AI and incorporates it, the AI Office has the powers of a market surveillance authority.²¹⁹ Where a deployer substantially modifies a general-purpose AI system for use in a high-risk AI context (i.e., fine-tunes a generative AI), it is considered a provider of a high-risk system and falls under the Act's high-risk regulation discussed at length above.²²⁰

Where a national market surveillance authority has reason for concern that a general-purpose AI system can be used in an *unmodified* way by deployers, and that it violates the requirements of the Act, the authority is directed to cooperate with the AI Office to ensure compliance. If the national market surveillance authority is stonewalled by the general-purpose AI provider, it is directed to submit a request to the AI Office to enforce its right to necessary oversight information.²²¹

C. AD HOC ELEMENTS AND OTHER STRANDS OF REGULATION

Finally, the Act contains ad hoc elements and other strands of regulation that fit neither into its core reliance on the NLF nor into its afterthought approach to general-purpose AI models. We conclude this Part by pointing to several of these, but again, this is not an exhaustive review. We discuss the Act's required disclosures to individuals; the Act's requirements of “Fundamental Rights Impact Assessments” and AI explanations to affected individuals; and strands of regulation regarding “innovation” that address regulatory sandboxes, real-world testing, and small businesses.

As mentioned above, the AI Act largely does not provide individual rights to affected persons, perhaps relying on rights established through other laws, such as the GDPR. However, Chapter IV of the Act (which consists of only one article, Article 50) establishes “Transparency obligations for providers and deployers of certain AI systems.”²²² These include the following: AI systems, such as chatbots, must be designed to inform people that they are interacting

218. *Id.* arts. 91–93.

219. *Id.* art. 75(1).

220. *Id.* art. 25(1)(c). This is actually true for any substantial modification of any type of AI system—anyone who does it becomes a provider—but it is particularly likely to come up in the case of fine-tuning generative AI.

221. *Id.* arts. 75(2)–(3).

222. *Id.* art. 50.

with an AI system.²²³ Deployers of an emotion recognition system must inform affected people that the system is in use.²²⁴ So must deployers of biometric categorization systems.²²⁵

Several of Article 50's transparency provisions address synthetic content, or "deep fakes." The providers of general-purpose AI systems that generate synthetic content must make sure outputs are marked as AI-generated in a machine-readable way.²²⁶ Deployers must disclose that images, audio, or video output has been artificially generated or manipulated.²²⁷ They also must disclose if text on matters of public concern has been artificially generated or manipulated.²²⁸

Two other provisions of the AI Act appear to be ad hoc responses to critiques made during drafting by EU's data privacy regulators, the European Data Protection Board (EDPB) and European Data Protection Supervisor (EDPS).²²⁹ These regulators criticized an earlier draft of the AI Act for failing to protect or really even address the rights of affected individuals.²³⁰ Consequently, the EDPB and EDPS urged the drafters to establish rights and remedies for "individuals subject to AI systems."²³¹ Specifically, they called for a "right to explanation" of AI decisions,²³² which Article 86 of the AI Act consequently provides.²³³ One of us has written detailed analysis of when and how this right applies, including how it interacts with similar rights established by the GDPR.²³⁴

Criticisms from the EDPB and EDPS also led to the adoption of risk mitigation at the deployer level: the Fundamental Rights Impact Assessment in Article 27. The EDPB and EDPS noted that initially, the AI Act (consistent

223. *Id.* art. 50(1) ("intended to interact directly with natural persons are designed and developed in such a way that the natural persons concerned are informed that they are interacting with an AI system").

224. *Id.* art. 50(3).

225. *Id.*

226. *Id.* art. 50(2).

227. *Id.* art. 50(4).

228. *Id.*

229. See EDPB-EDPS Joint Opinion, *supra* note 106.

230. *Id.* ¶ 18 ("Whether they are end-users, simply data subjects or other persons concerned by the AI system, the absence of any reference in the text to the individual affected by the AI system appears as a blind spot in the Proposal.").

231. *Id.* ("the EDPB and the EDPS urge the legislators to explicitly address in the Proposal the rights and remedies available to individuals subject to AI systems").

232. *Id.* ¶ 60 ("A right to explanation should provide for additional transparency.").

233. AI Act, *supra* note 2, art. 86.

234. Margot E. Kaminski & Gianclaudio Malgieri, *The Right to Explanation in the AI Act, in THE EU ARTIFICIAL INTELLIGENCE ACT: A THEMATIC COMMENTARY* (Gianclaudio Malgieri, Gloria González Fuster, Alessandro Mantelero & Gabriela Zanfir-Fortuna eds., forthcoming 2026).

with the NLF) placed risk mitigation requirements upon providers only, and not upon deployers.²³⁵ They explained that sometimes AI system users (deployers) should be the ones conducting risk mitigation.²³⁶ Consequently, the AI Act's drafters added Article 27. It requires deployers that are government entities, private entities providing public services, and deployers of high-risk AI systems that price insurance or determine creditworthiness, to conduct a Fundamental Rights Impact Assessment (FRIA) prior to first use.²³⁷ The FRIA must among other things, identify the categories of persons and groups likely to be affected and the specific risks of harm to them, and outline measures taken in case risks materialize, including "internal governance and complaint mechanisms."²³⁸ The results must be reported to the market surveillance authority.²³⁹

Finally, the AI Act contains numerous provisions that it describes as "measures in support of innovation."²⁴⁰ We address these as separate strands of regulation, distinct from the NLF framework. They are not strictly speaking ad hoc, as they have been part of the Act's framework since early drafting. These strands include provisions establishing AI regulatory sandboxes, provisions on testing systems in real-world conditions, and provisions regarding the regulation of small businesses.

A regulatory sandbox is a type of temporary and experimental regulatory regime that might provide a break from regulation in exchange for supervision.²⁴¹ Typically, a sandbox entails public-private collaboration, coupled with clarifying guidance from a regulator. They were first used in the context of U.K. Fintech regulation.²⁴² As contemplated in the AI Act, regulatory sandboxes aim to foster innovation, support startups, improve legal certainty, contribute to the sharing of best practices, and contribute to "evidence-based regulatory learning."²⁴³

235. See EDPB-EDPS Joint Opinion, *supra* note 106, ¶¶ 20–21.

236. *Id.*

237. AI Act, *supra* note 2, art. 86(1); see also *id.* Annexes III(5)(b)–(c) ("AI systems intended to be used to evaluate the creditworthiness of natural persons or establish their credit score, with the exception of AI systems used for the purpose of detecting financial fraud; AI systems intended to be used for risk assessment and pricing in relation to natural persons in the case of life and health insurance").

238. *Id.* arts. 86(1)(a)–(f).

239. *Id.* art. 86(3).

240. See *id.* ch. VI (titled "Measures in Support of Innovation"); *id.* art. 57 (titled "AI Regulatory Sandboxes").

241. See generally Sofia Ranchordás & Bart van Klink, *Special Issue Experimental Legislation in Times of Crisis*, 11 L. & METHOD 1 (2022).

242. *Id.* at 6.

243. AI Act, *supra* note 2, art. 57(9).

The AI Act requires that member states establish at least one AI regulatory sandbox by August 2026.²⁴⁴ AI regulatory sandboxes are described as a “controlled environment . . . [that] facilitates the development, training, testing and validation of innovative AI systems for a limited time before their being placed on the market or put into service pursuant to a specific sandbox plan agreed between the providers . . . and the competent authority.”²⁴⁵ An AI provider that has been subject to a sandbox may use its sandbox “exit report” to later demonstrate compliance with the AI Act.²⁴⁶

An overlapping strand of regulation has to do with testing. Multiple provisions of the Act address testing AI systems in real-world conditions. Real-world testing is contemplated within regulatory sandboxes.²⁴⁷ Article 60 additionally governs real-world testing of up to a year outside of the sandbox context, subject to a real-world testing plan that is submitted to the market surveillance authority.²⁴⁸ Article 61 requires informed consent from test subjects.²⁴⁹

The AI Act also contains several measures that affect its application to small businesses and startups.²⁵⁰ For example, it requires that small and medium-sized businesses be given priority access to regulatory sandboxes.²⁵¹ It also requires that member states facilitate participation by small businesses in the standardization process.²⁵² It requires that small businesses and startups be charged reduced fees for conformity assessments.²⁵³ And the Act requires the Commission to develop guidelines for simplified compliance with the quality management system (required under Article 17) for businesses with fewer than ten employees generating under €2 million in revenue.²⁵⁴

IV. ANALYSIS

In this last Part, we offer some analysis. We begin with a discussion of whether the AI Act is better understood as an instantiation of the

244. *Id.* art. 57(1).

245. *Id.* art. 57(5).

246. *Id.* art. 57(7).

247. *Id.* art. 57(5).

248. *Id.* arts. 60(1), 60(4).

249. *Id.* art. 61(1) (“freely-given informed consent shall be obtained from the subjects of testing prior to their participation in such testing”).

250. *Id.* art. 62.

251. *Id.* art. 62(1)(a).

252. *Id.* art. 62(1)(d).

253. *Id.* art. 62(2).

254. *See* Commission Recommendation, art. 2(3), 2003 O.J. (L 124) 39 (defining “microenterprise”), <https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2003:124:0036:0041:EN:PDF>.

precautionary principle, or a law promoting the uptake of AI. We then discuss how the Act constructs high-risk AI systems through a product-safety framing, and the consequences of doing so. We emphasize that the Act should be understood as the unique product of its drafting story. We describe it as a “legal exoskeleton”: a hard-law framework with a softer-law interior that leaves many open questions. We close with some musings on global power politics and the future of the law and its global influence.

A. PRECAUTION, OR A BID FOR AI BUSINESS?

If there's one concept American lawyers tend to think of with respect to the difference between EU and U.S. regulators, it's the precautionary principle, which emphasizes caution about new technologies. In the United States, we innovate first, ask questions later, while in Europe the red tape is layered on thick—or so the story goes.²⁵⁵

For this reason, it is notable that the AI Act puts the “uptake of human-centric and trustworthy artificial intelligence” on an equal purposive footing to the “high level of protection of health, safety, fundamental rights” it aims to “ensure.”²⁵⁶ The Act centrally prioritizes the uptake of AI. One way to understand the AI Act, against the typical backdrop of EU regulation, is that the EU decided to attract more AI use.²⁵⁷ Being the first in the world to pass a comprehensive AI law may have been a play to capture more of the global market for AI, with protection of EU citizens a secondary goal.²⁵⁸

This would certainly help explain the reliance on the product-safety framing and the shunting of fundamental rights concerns to a framework of self-certification. As one of us has argued elsewhere, there is a growing global convergence on risk regulation as governance of AI.²⁵⁹ Risk regulation is

255. Douglas A. Kysar, *It Might Have Been: Risk, Precaution and Opportunity Costs*, CORN. L. FAC. PUBL'NS, PAPER 50, 1, 3–4 (2006) (noting United States favors cost-benefit analysis that predicts, weighs, and aggregates consequences of policy proposals to identify “welfare-maximizing uses of public resources,” while EU approach to risk regulation is associated with precautionary principle). *But see* Jonathan B. Wiener, *Whose Precaution After All? A Comment on the Comparison and Evolution of Risk Regulatory Systems*, 13 DUKE J. COMPAR. & INT'L L. 207, 213–15 (2003) (noting the common perception that the EU favors precaution and the United States favors more permissive regulation is oversimplified).

256. AI Act, *supra* note 2, art. 1(1).

257. *See* Troels Krarup & Maja Horst, *European Artificial Intelligence Policy as Digital Single Market Making*, 10 BIG DATA & SOC'Y 1 (2023).

258. *See, e.g.,* Daniel Leufer, Fanny Hidvegi & Alessia Zornetta, *The Pitfalls of the European Union's Risk-Based Approach to Digital Rulemaking*, 71 UCLA L. REV. DISCOURSE 156, 166 (2024) (“[T]he Commission's strategy on AI since 2018 has focused primarily on boosting AI uptake across the EU, with measures to tackle the negative impacts of such uptake coming as a secondary consideration.”).

259. Kaminski, *supra* note 51, at 1396.

typically “ex ante, systemic, and concerned with aggregate outcomes,”²⁶⁰ rather than focused on ex post accountability for individualized harms. Among risk regulation’s “policy baggage” is that risk regulation often “takes as its starting point that a technology must be fixed so that it can be used.”²⁶¹ Thus, not only does the Act expressly state that it wants to increase uptake of AI, but by adopting a risk regulation framework, it essentially takes a stance that AI is fixable, but inevitable.

While risk regulation is not always incompatible with a precautionary approach—precaution and risk analysis often go hand-in-hand, after all—the legal and intellectual framework of this particular approach to risk regulation contrasts with the possibility of stronger prohibitions, a robust system of individual redress, and ex post liability for harms. While the AI Act does notably contain bans, they are narrow and limited. And the choice to create a law under the NLF framework is a break from what one might typically expect from an EU concerned with vindicating individual fundamental rights.

B. LEGALLY CONSTRUCTING HIGH-RISK AI SYSTEMS THROUGH PRODUCT SAFETY

The AI Act constructs the legal concept of “high-risk AI systems” through the values and institutions of product safety regulation.²⁶² It treats AI systems as though they are elevators or children’s toys, rather than bureaucratic systems for decision-making.²⁶³ Not only does the Act characterize AI systems as products, it envisions AI systems as though they are products designed for and used for a particular intended purpose. Inherently, general-purpose AI systems do not readily fit into this regulatory model.²⁶⁴

In characterizing AI systems as products, the Act imports the values of product safety regulation. It prioritizes protection of health and safety over the protection of fundamental rights.²⁶⁵ It relies on ex post measures such as product recalls, which indicates that some level of harm is acceptable, up to a point. This is not how the EU generally treats fundamental rights. Typically,

260. *Id.* at 1369.

261. *Id.* at 1397.

262. For a discussion of the method of legal construction of technology, see Margot E. Kaminski & Meg Leta Jones, *Constructing AI Speech*, 133 YALE L.J. 1212 (2024). Here, we discuss the objects (AI as product), values (the values of product safety regulation), and institutions (the institutions of the NLF). *See also* Petit & Almada, *supra* note 57.

263. *See also* Veale & Zuiderveen Borgesius, *supra* note 57, at 102.

264. *See generally* Claire Boine & David Rolnick, *Why the AI Act Fails to Understand Generative AI*, 26 MINN. J. L., SCI. & TECH. 61 (2025).

265. Recall that AI systems with safety implications undergo third-party conformity assessments, while AI systems implicating fundamental rights get self-certification.

EU law turns on whether a right has been violated, not whether a certain level of measurable harm has been caused.²⁶⁶

The AI Act constructs AI systems as particular objects (products), through particular institutions (product safety regulators). The AI Act piggybacks on the existing institutional framework of the NLF (albeit as discussed with some important caveats). Its infrastructure relies on the existence of two classes of member state institutions: (a) market surveillance authorities, responsible for market surveillance and ensuring compliance, and (b) notifying authorities, responsible for designating third-party conformity assessment bodies. Recall that member states may designate existing NLF institutions for these roles. This matters because the Act will be primarily enforced by institutions trained in operationalizing product safety values rather than fundamental rights values.²⁶⁷ For example, in Finland, the Transport and Communications Agency, a preexisting market surveillance authority, is responsible for the Act's enforcement.²⁶⁸

Member states do have discretion in deviating from these institutions, however. For example, Italy instead designated its national cybersecurity agency as its market surveillance authority.²⁶⁹ Other member states have designated their data protection authority as the market surveillance authority.²⁷⁰ Still others have created new bodies specifically for regulating AI systems.²⁷¹ The array of different institutions with different degrees and kinds of expertise, resources, and values will greatly affect how the Act is implemented in practice. In short, while some of these institutions will bring values other than products safety values to the enforcement table, there is no guarantee these institutions will be well-versed in the CJEU's fundamental

266. See Mireille Hildebrandt, *Beyond the GDPR?*, Presentation for the COHUBICOL ERC ADG Project 1, 18 (2023), <https://www.cohubicol.com/assets/uploads/response-hildebrandt-purtova.pdf>. See also Katerina Demetzou, *Data Protection Impact Assessment (2014–24)* (PhD dissertation, Radboud Business Law Institute) (discussing risks to rights under the GDPR).

267. See Veale & Zuiderveen Borgesius, *supra* note 57, at 112 (“The enforcement mechanism is a creature of product safety.”).

268. *Overview of All AI Act National Implementation Plans*, EU ARTIFICIAL INTELLIGENCE ACT (Nov. 8, 2024), <https://artificialintelligenceact.eu/national-implementation-plans/> (“A draft implementing act from October, 2024, appoints 10 already existing market surveillance authorities The Finnish Transport and Communications Agency will act as the single point of contact.”).

269. *Id.* (“A legislative proposal from May 2024 designates the National Cybersecurity Agency (Agenzia per la Cybersicurezza Nazionale, ACN) as market surveillance authority with monitoring, inspection and enforcement powers in relation to AI systems”).

270. *Id.* See, e.g., Luxembourg & Malta.

271. *Overview of All AI Act*, *supra* note 268 (Poland and Romania sections).

rights decisions. This great variety of institutions will also lead to a great deal of heterogeneity in implementation across member states.

The AI Act does not stand in isolation, however, and it is important to be aware of what does and does not exist around it. For example, product safety regulation typically interacts with product liability, with liability providing a backstop to regulation and an opportunity for individual redress. However, there is no corresponding harmonized product liability law on the EU level. The proposed AI Liability Directive, which would have done some harmonizing on AI liability, is now dead.²⁷² Member states' approaches to liability accordingly will not be harmonized with respect to AI.

Without a backstop product liability directive, the Act's central reliance on risk regulation takes on more significance. That reliance on risk regulation affects both substance and institutional design. It results in a strange fit between the tools the Act uses and the harms it aims to prevent.²⁷³ The AI Act turns on a quantified definition of "risk," where human rights violations are typically not quantifiable.²⁷⁴ With only a few exceptions,²⁷⁵ the Act does not afford specific individual rights. For a law purportedly aimed in no small part towards fundamental rights protection, this is bizarre.²⁷⁶

From an institutional design perspective, the lack of individual rights or redress is significant. It potentially makes it significantly harder to get cases on fundamental rights protections under the Act before the Court (the CJEU). A contrast with the GDPR here may be helpful. As outlined in Part I, typically, a case comes before the CJEU when it is referred by a member state's national court on a question of EU law. Under the GDPR, an individual whose rights are violated can lodge a complaint with a member state's data protection

272. See Caitlin Andrews, *European Commission Withdraws AI Liability Directive from Consideration*, INT'L ASS'N PRIV. PROS. (Feb. 12, 2025), <https://iapp.org/news/a/european-commission-withdraws-ai-liability-directive-from-consideration> (describing the withdrawal of the proposed directive in Feb. 2025); Cynthia Kroet, *Lawmakers Reject Commission Decision to Scrap Planned AI Liability Rules*, EURO NEWS (Feb. 18, 2025), <https://www.euronews.com/next/2025/02/18/lawmakers-reject-commission-decision-to-scrap-planned-ai-liability-rules> (showing attempts to revive the directive). But see Deimante Rimkute, *AI Liability After the AILD Withdrawal: Why EU Law Still Matters?*, OXFORD BUS. L. BLOG (Apr. 1, 2025), <https://blogs.law.ox.ac.uk/oblb/blog-post/2025/04/ai-liability-after-aild-withdrawal-why-eu-law-still-matters> (arguing that existing EU law on products liability nonetheless pushes towards harmonization).

273. Kaminski, *Regulating the Risks of AI*, *supra* note 51, at 1400.

274. *Id.* at 1401 ("A second, central problem of AI risk regulation is that the risks raised by AI systems are varied, not always quantifiable, often contested, and sometimes excruciatingly or even impossibly hard to define.").

275. E.g., AI Act, *supra* note 2, art. 86, and some notification requirements, *id.* art. 50.

276. See EDPB-EDPS Joint Opinion, *supra* note 106, ¶ 18 ("Whether they are end-users, simply data subjects or other persons concerned by the AI system, the absence of any reference in the text to the individual affected by the AI system appears as a blind spot in the Proposal.").

authority, can sue a member state's data protection authority (DPA), or can sue a data controller or processor in court.²⁷⁷ If that case goes up before the national court, that court can then refer questions on data protection law to the CJEU. (This process itself takes a long time!) This has in practice enabled the CJEU to act as the institutional rights-protective backstop to the regulatory regime outlined in the GDPR.²⁷⁸

The AI Act, by contrast, lacks as robust of a fundamental rights backstop. On the one hand, the Charter is still the source of fundamental rights and backstops the Act whether it is extensively operationalized or not.²⁷⁹ On the other, the Act for the most part channels enforcement into market surveillance authorities and unlike the GDPR does not establish a right to sue those authorities, or to sue AI providers.²⁸⁰ Whether or not it is possible to sue designated market surveillance authorities is a matter of member state law. This makes it potentially harder for the CJEU to serve as a rights-protective backstop to the AI Act's regime.²⁸¹

277. See GDPR, *supra* note 19, arts. 77–79.

278. See Margot E. Kaminski & Meg Leta Jones, *American's Guide to the GDPR*, 98 DENV. L. REV. 93 (2021).

279. See Simona Demková, *The EU's Artificial Intelligence Laboratory and Fundamental Rights*, in REDRESSING FUNDAMENTAL RIGHTS VIOLATIONS BY THE EU 391, 411 (Melanie Fink ed., 2024) (“the AI Act will need to be applied in conjunction with the existing EU law, including the rules on remedies and existing data protection rules”).

280. *Id.* at 409 (“it remains to be stressed that judicial remedies are rather limited in the context of AI-powered conduct based on composite administrative procedures involving actors at EU and Member State levels”); *id.* at 415–16 (“Section 4 in the final version of the Act . . . provides however only provides a limited consolidation of the calls for enhancing access to justice against the risks of AI . . . the remedies under the AI Act are essentially two-fold: (a) a product-related complaint mechanisms before the designated market surveillance authorities; (b) the right to an explanation of individual decision-making when the latter is made on the basis of a high-risk AI output”).

281. Important caveat: with general-purpose AI, the Court will be able to hear questions about the Commission. See *Court of Justice of the European Union (CJEU): Overview*, EU, https://european-union.europa.eu/institutions-law-budget/institutions-and-bodies/search-all-eu-institutions-and-bodies/court-justice-european-union-cjeu_en (“ensuring the EU takes action . . . sanctioning EU institutions”).

It is possible, however, that some AI Act questions will make it to the CJEU regardless,²⁸² including some provisions with direct effect,²⁸³ or where individuals have rights to sue under other legal regimes.

Once again, however, the AI Act is not the only law governing AI in the EU. The Act depends on the existence of other laws that afford fundamental rights protections, consumer protections, and other requirements. These are the legal waters AI providers swim in. The Charter continues to provide protections for fundamental rights. The GDPR continues to afford individuals data protection rights. The Digital Services Act establishes specific obligations

282. Art. 86 may also come up as the one explicit individual remedy provided in the AI Act. *See, e.g.*, Request for a Preliminary Ruling from the Sofiyski Rayonen Sad (Bulgaria), 2025 O.J. (C-806/1080), <https://eur-lex.europa.eu/eli/C/2025/1080/oj/eng> (referring multiple questions about the AI Act to the CJEU in a case involving the Consumer Protection Directives) (for example: “Must Article 86(1) of Regulation (EU) 2024/1689 (1) be interpreted as meaning that the consumer has the right, within the meaning of Directives 2011/83/EU (2) and 93/13/EEC, (3) to know from the service provider how and with the aid of what elements [and] parameters automated decisions (invoices) were generated on the basis of data which the trader collected automatically in the context of a contract for the provision of mobile telecommunications services? . . . 4. [Must Art. 86] be interpreted as permitting the court to demand from the trader the black box data, the source code and the algorithm relating to the way in which automated decisions are made under the consumer contract? 5. Must Article 86(1) of Regulation (EU) 2024/1689, read in conjunction with Article 47 of the Charter of Fundamental Rights of the European Union, read in conjunction with Article 38 of the Charter, and with Directive 2011/83/EU, be interpreted as meaning that an automated decision generated by a trader under a contract with a consumer for mobile telecommunications services permits that automated decision to be reviewed by a human being, a judge, during real judicial proceedings? Must those provisions be interpreted as meaning that automated decisions . . . are subject to human review by a judge in real judicial proceedings? . . . 6. Must recitals 7 and 8 and Article 95(2)(a) of Regulation (EU) 2024/1689—the AI Act—and Directive 2011/83/EU . . . be interpreted as meaning that, where an automated decision-making system is operated and used . . . in the consumer contract, lawyers or senior judicial officers . . . with high moral and ethical standards must be involved in order to guarantee a transparent, effective and human-centric information system which takes account of fundamental rights? . . . 10. Must Article 5(1) of Directive 93/13/EEC and Article 86(1) of Regulation (EU) 2024/1689 be interpreted as meaning that the automatically generated invoices arising from a consumer contract . . . must be written in plain, intelligible language and the consumer has the right to demand an explanation from the trader as to how and by what algorithm the decision was made?”).

283. *Direct Effect*, EUR. INDUS. RELS. DICTIONARY (Feb. 15, 2017), <https://www.eurofound.europa.eu/en/european-industrial-relations-dictionary/direct-effect> (“the CJEU identified three situations necessary to establish the direct effect of primary EU law. These are that: the provision must be sufficiently clear and precisely stated; it must be unconditional and not dependent on any other legal provision; it must confer a specific right upon which a citizen can base a claim.”); *see also The Direct Effect of European Union Law*, EUR-LEX, <https://eur-lex.europa.eu/EN/legal-content/summary/the-direct-effect-of-european-union-law.html> (“in line with the general principles, this applies only under the condition that the rules are sufficiently clear, precise and relevant to the situation of the individual litigant (direct effect as clarified by the *Politi v Ministero delle finanze* Court judgement)”).

for large online platforms. The Copyright in the Digital Single Market Directive addresses copyright concerns.²⁸⁴

There are two important consequences of the Act's reliance on other EU regulations as backdrop. First, the Act's reliance on existing fundamental rights protections in part justifies its product safety approach.²⁸⁵ The Act purportedly can rely on risk regulation as its central mechanism precisely because individuals are afforded rights and redress through other regulations.

A second consequence of the Act's reliance on other EU law, however, is that if you have substantive questions about certain aspects of AI regulation, your answers may lie elsewhere. This certainly makes things trickier for American lawyers trying to assess the legality of a particular AI system. For example, if you want to determine the legality of a particular biometrics system, you will have to look both to the AI Act and to the GDPR.²⁸⁶ If you want to determine how an online platform can use AI in content moderation, you will have to look to both the AI Act and the DSA. If you want to understand the EU's approach to copyright and training data, you will have to look to both the AI Act and EU copyright law.²⁸⁷

The relationship to data protection in particular is fraught. Both the AI Act and the GDPR are concerned with data, but towards fundamentally different ends. The GDPR at its core protects the “data subject”—individuals affected by the collection, processing, and use of personal data, protected not just under the GDPR but also under the Charter. The AI Act has no such concerns.²⁸⁸ It's more concerned with mitigating inaccuracy and ensuring fit to purpose. However, the AI Act does not displace data protection law;²⁸⁹ nor does it displace data protection institutions, on which it occasionally in fact relies.²⁹⁰

284. Directive (EU) 2019/790, of the European Parliament and of the Council (on copyright and related rights in the Digital Single Market), 2019 O.J. (L130) 92; *see also* Quintais, *supra* note 41.

285. *See* Eike Graef & Paul Nemitz, *Addressing the Challenge of Protecting Fundamental Rights Through AI Regulation in the European Union*, 71 UCLA L. REV. DISCOURSE 144, 151 (2024) (“It is important to look at the AI Act proposal together with these other elements, because the different initiatives and laws are designed to complement and strengthen each other.”).

286. *See, e.g.,* Demková, *supra* note 279.

287. *See* Quintais, *supra* note 41.

288. *See* EDPB-EDPS Joint Opinion, *supra* note 106.

289. With the exceptions discussed above, in AI Act, *supra* note 2, arts. 10(5)(b), 59.

290. *See, e.g., id.*, art. 70(9) (establishing the European Data Protection Supervisor as the market surveillance authority for EU providers).

C. THE ACT AS PRODUCT OF ITS DRAFTING STORY

The AI Act is the product of its messy drafting story. The Act may have started as a transplant of the NLF to high-risk AI systems, which itself is complex enough. But it ended up as a complicated Frankenstein.

The Act's approach to general-purpose AI was added after-the-fact. Several of its provisions on fundamental rights were added to respond to data protection regulators and other critics. And its bans on certain AI uses can almost all be traced to particular news items about particular applications of predictive AI. To understand the AI Act, then, you need to both be aware of its core reliance on the NLF, and be aware that as a law, it is quintessentially the product of its times.

It's also worth noting that while the AI Act took three years to draft,²⁹¹ many other EU regulations have had a far longer runway. Contrast, for example, the GDPR. European member states have had data protection law for decades, some since the 1970s. The EU-wide Data Protection Directive went into effect in the mid-1990s. The EU Charter, established in 2000, contains a fundamental right to data protection. By the time the GDPR was promulgated, there was both buy-in to and considerable infrastructure for the project of European data protection law. A similar point can be made about the slower progression in online platform regulation from the 2000 E-Commerce Directive to the 2022 Digital Services Act. By contrast, the AI Act was established as a regulation, over the course of only three years, with no preceding directive creating buy-in from citizens or member states. Contrasted with the EU's data protection regime and online platform regulation, it does not reflect a similar sort of bubble-up practical consensus developed over time. As a result, the AI Act may face legitimacy problems, leading to future amendments and/or difficulties with compliance.

D. THE ACT AS LEGAL EXOSKELETON

We describe the AI Act at its core as a *legal exoskeleton*: a hard-law legal framework surrounding a softer-law interior. The AI Act's central reliance on delegating—on having somebody else fill in most of the substantive blanks—means that at its core, it could turn out to be quite permissive. The Act delegates much of its substance to multiple potential actors, including technical standards-setting bodies and the Commission. Each bring with them democratic legitimacy problems, and a not insignificant likelihood of producing outcomes reflecting power politics. We here discuss both the delegation to technical standards-setting bodies, and the more recent

291. *Historic Timeline*, EU ARTIFICIAL INTELLIGENCE ACT, <https://artificialintelligenceact.eu/developments/>.

development of the Commission's General-Purpose AI (GPAI) Code of Practice.

A law can be softer or harder along different axes.²⁹² These include how mandatory versus voluntary a law is, who decides the rules, whether the rules are vague or specific, and more. The AI Act, like a number of recent European regulations, contains a complex and deliberate mixture of both harder and softer law. Its status as a regulation, rather than a directive, makes it directly binding on member states and thus harder law than a directive. Its enforcement mechanisms, which include large fines, also evidence a serious degree of hardness. However, the Act's central reliance on fuzzier language, instantiated in standards to be developed by technical standards-setting organizations, indicates a softer belly.

While the Act establishes a large and complicated accountability framework, we still don't know most of the actual substantive requirements for AI systems. The AI Act requires things like "the estimation and evaluation of the risks that may emerge"²⁹³ and that "[t]raining, validation and testing data sets shall be relevant, sufficiently representative, and to the best extent possible, free of errors."²⁹⁴ Both of these requirements, among many others, leave plenty of room for doubt about whether a particular provider is in compliance or not.

The main thing that providers of AI systems will want is certainty. Thus the rapid move to technical standards-setting. As soon as the joint technical committee of CEN/CENELEC finishes developing its AI Act standards, that document will probably become the de facto legal regime. Even though the standards that CEN/CENELEC come up with are not themselves technically binding law, they nonetheless effectively will have the force of law because compliance with them will establish a presumption of conformity with the law.

Until those standards are released, we do not know how substantive, detailed, or technical they will actually be. They may provide specific metrics or benchmarks, or they may echo the Act's existing fuzziness, or they may do both. At its core, then, the Act could end up with (a) less-than rigorous requirements negotiated in large part by private actors, (b) "technical" requirements that remain high-level and more like standards than like rules, or likely (c) some combination of the two.

292. See Kenneth W. Abbott & Duncan Snidal, *Hard and Soft Law in International Governance*, 54 INT'L ORG. 421, 421 (2000).

293. AI Act, *supra* note 2, art. 9(3).

294. *Id.* art. 10(3).

By delegating central decision-making, the Act at its core has democratic legitimacy problems.²⁹⁵ One set of commentators has noted that by putting so central of an emphasis on technical standards-setting to govern fundamental rights, the AI Act puts a “constitutional bomb” under the NLF, a framework that has otherwise worked well to govern product safety in the EU.²⁹⁶ As odd as it may be to rely on market surveillance authorities for oversight of fundamental rights when they lack such expertise, it is downright bizarre to delegate fundamental-rights decision-making to private standards-setting organizations.²⁹⁷

On the other hand, the EU’s standards-setting organizations are not purely private, as they are in the United States. The EU’s formal process for requesting technical standards and establishing them already involves more public sector decision-making and oversight than U.S. incorporation of private standards.²⁹⁸ The CJEU has case law establishing that standards that are given legal effect are to be treated more like actual law.²⁹⁹ For example, the CJEU has required that technical standards under copyright law be publicly accessible because they have the force of law.³⁰⁰ Moreover, the AI Act imposes certain participation requirements on the AI standards-setting process in particular.³⁰¹

295. See Marta Cantero Gamito & Christopher T. Marsden, *Artificial Intelligence Co-Regulation? The Role of Standards in the EU AI Act*, 32 INT’L J. L. & INFO. TECH. 1 (2024).

296. Veale & Zuiderveen Borgesius, *supra* note 57, at 105.

297. One commentator, speaking of parallel processes at NIST in the United States, has called these kinds of AI ethical governance documents crafted through technical standards-setting organizations “un-standards.” Bryan H. Choi, *NIST’s Software Un-Standards*, 9 GEO. L. TECH. REV. 65 (2025).

298. See Emily S. Bremer, *American and European Perspectives on Private Standards in Public Law*, 91 TULANE L. REV. 325 (2016).

299. Case C-588/21, Public.Resource.Org, Inc. and Right to Know CLG v Eur. Comm’n, ECLI:EU:C:2024:201, ¶ 80 (Mar. 2024) (“In the light of the foregoing considerations, it must be held, in accordance with the case-law referred to in paragraph 70 of the present judgment, that the requested harmonised standards form part of EU law.”).

300. *Id.* ¶ 85 (“In those circumstances, it must be held that there is an overriding public interest, within the meaning of the last clause of Article 4(2) of Regulation No 1049/2001, justifying the disclosure of the requested harmonised standards.”).

301. See, e.g., AI Act, *supra* note 2, art. 40(3) (“The participants in the standardisation process shall seek to promote investment and innovation in AI, including through increasing legal certainty, as well as the competitiveness and growth of the Union market, to contribute to strengthening global cooperation on standardisation and taking into account existing international standards in the field of AI that are consistent with Union values, fundamental rights and interests, and to enhance multi-stakeholder governance ensuring a balanced representation of interests and the effective participation of all relevant stakeholders in accordance with Articles 5, 6, and 7 of Regulation (EU) No 1025/2012.”); *id.* art. 62(d); *id.* recital 121 (“A balanced representation of interests involving all relevant stakeholders in the development of standards, in particular SMEs, consumer organisations and environmental and social stakeholders in accordance with Articles 5 and 6 of Regulation (EU) No 1025/2012 should therefore be encouraged”); *id.* recital 143.

The Act contemplates that the Commission can reject standards that it finds do not comport with the law and instead craft its own “common specifications.”³⁰² That is, it backstops softer law with a public law option. (However, Commission-made law may itself be problematic—more on this in a moment.)

Understanding the AI Act as legal exoskeleton brings us to a crucial observation: the Act's central reliance on technical standards makes it vulnerable to international realpolitik. On the one hand, the EU passed the AI Act ostensibly to develop AI with built-in European values. On the other, by centrally relying on technical standards, the Act opens a side door to international influence. Other standards-setting organizations have already promulgated AI standards, or are far along in the process. China has recently increased its participation in international standards-setting.³⁰³ In the United States, NIST issued its AI risk management framework in 2023,³⁰⁴ and has turned its attention specifically to influence global standards-setting.³⁰⁵

Whatever standards the European Standardization Organizations issue will have an eye to international consensus-building. (In fact, they have to, under the WTO, to the extent anybody still follows the rules of international trade.³⁰⁶)

302. AI Act, *supra* note 2, art. 41(1)(a)(iii) (when “(iii) the relevant harmonised standards insufficiently address fundamental rights concerns”); *see also id.* art. 40(2).

303. *See generally* Marta Cantero Gamito, *The Influence of China in AI Governance Through Standardisation*, 47 TELECOMM. POL'Y 102673 (2023).

304. *AI Risk Management Framework*, NAT'L INST. OF STANDARDS & TECH., <https://www.nist.gov/itl/ai-risk-management-framework>.

305. Jesse Dunietz, Elham Tabassi, Mark Latonero & Kamie Roberts, *A Plan for Global Engagement on AI Standards*, NAT'L INST. OF STANDARDS & TECH. (July 26, 2024), <https://www.nist.gov/publications/plan-global-engagement-ai-standards>; *see also* *Winning the Race: America's AI Action Plan*, WHITE HOUSE (July 2025), <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf> (Pillar III: Lead in International AI Diplomacy and Security, including “Counter Chinese Influence in International Governance Bodies”: “leverage the U.S. position in international diplomatic and standard-setting bodies to vigorously advocate for international AI governance approaches that promote innovation, reflect American values, and counter authoritarian influence”).

306. *See* World Trade Organization, Agreement on Technical Barriers to Trade (TBT), ¶ 2.2, Apr. 15, 1995, https://www.wto.org/english/docs_e/legal_e/17-tbt.pdf (“Members shall ensure that technical regulations are not prepared, adopted or applied with a view to or with the effect of creating unnecessary obstacles to international trade.”); *see also id.* ¶ 2.4 (“Where technical regulations are required and relevant international standards exist or their completion is imminent, Members shall use them, or the relevant parts of them, as a basis for their technical regulations except when such international standards or relevant parts would be an ineffective or inappropriate means for the fulfilment of the legitimate objectives pursued . . .”). *See generally* *Technical Information on Technical Barriers to Trade*, WORLD TRADE ORG., https://www.wto.org/English/tratop_e/tbt_e/tbt_info_e.htm.

The AI Act itself notes this dynamic.³⁰⁷ There are signs that the Commission is aware of this risk themselves; they purportedly left out one of the European Standardization Organizations out of fear about U.S. and Chinese influences on the process.³⁰⁸

There are first-mover advantages in the standards-setting race. The EU may have been the first to pass omnibus AI hard law, but it was not the first to promulgate standards.³⁰⁹ This means that it is as likely to import U.S. and Chinese values via standards-setting as it is to effectively establish European AI standards for export. Again, the Commission's ability to reject and thus check delivered standards may be crucial in determining which way the influence flows. However, unlike the Court, the Commission is not a human rights body. The Commission is less incentivized than, e.g., the Court, to push back on inadequate standards for fundamental rights reasons.³¹⁰

We end with a related recent plot twist concerning the recently developed GPAI Code of Practice. Again, there are many open questions about how this will all play out in the longer run. But the recent focus on the Code of Practice suggests a sort of shell game in where the Act's soft-law lawmaking may be taking place—and that the Commission is also not immune from political pressures.

Article 56 of the Act establishes a policy-making-qua-convening role for the AI Office at the Commission, with respect to general-purpose AI

307. See AI Act, *supra* note 2, art. 40(3) (“to contribute to strengthening global cooperation on standardisation and taking into account existing international standards in the field of AI that are consistent with Union values, fundamental rights and interests”).

308. Luca Bertuzzi, *Commission Leaves European Standardisation Body Out of AI Standard-Setting*, EURACTIV (Dec. 7, 2022), <https://www.euractiv.com/section/tech/news/commission-leaves-european-standardisation-body-out-of-ai-standard-setting/> (“[T]he European Commission set out its strategy to become more assertive in the way it participated in standard-setting, where it considered that non-European companies, particularly American and Chinese, have gained the upper hand. The strategy came as a slap in the face to ETSI, which the Commission accused of being held hostage by non-European influences and requested internal reform to give more weight to national standardisation bodies.”); see also AI Act, *supra* note 2, art. 40(3) (“taking into account existing international standards in the field of AI that are consistent with Union values, fundamental rights and interests”).

309. Technically, the Colorado AI Act passed first, but it is primarily an antidiscrimination law, not as omnibus.

310. See, e.g., the saga of *Schrems I* and *Schrems II*, in which the Court twice found that the Commission's negotiated Safe Harbor agreement with the U.S. government violated fundamental rights under the Charter. Case C-362/14, Maximilian Schrems v. Data Protection Commissioner, ECLI:EU:C:2015:650 (Oct. 6, 2015) [*Schrems I*]; Case C-311/18, Data Protection Commissioner v. Facebook Ireland Ltd., ECLI:EU:C:2020:559 (July 16, 2020) [*Schrems II*].

models.³¹¹ It tasks the AI Office with “encourag[ing] and facilitat[ing] the drawing up of codes of practice at Union level in order to contribute to the proper application of this Regulation, taking into account international approaches.”³¹² Consequently, the AI Office convened a group of experts to come up with the draft GPAI Code of Practice,³¹³ which was subsequently ratified by the Commission.³¹⁴

Compliance with the GPAI Code of Practice serves to demonstrate compliance with the Act, at least until the standards-setting organizations arrive at a harmonized standard for general-purpose AI.³¹⁵ In fact, it appears that a general-purpose AI model provider could choose its own legal adventure and comply with the GPAI Code of Practice *instead of* any harmonized standard.³¹⁶ So if the GPAI Code of Practice is comparatively weak, while the technical-standards-setting output is more rigorous, the Code could provide a path of least regulatory resistance.

Unsurprisingly, the availability of this option led to “intense lobbying.”³¹⁷ The Code of Practice was drafted involving over one thousand stakeholders from academia, civil society, and industry. But GPAI companies had a “special seat at the table,” including privileged access to the latest version of the text.³¹⁸ Nonetheless, multiple U.S. companies indicated they would not sign on to the Code.³¹⁹ Ultimately, the Commission withstood the pressure, and both the

311. AI Act, *supra* note 2, art. 56(2) (“at least the obligations provided for in Articles 53 and 55”).

312. *Id.* art. 56(1).

313. *The General-Purpose AI Code of Practice*, EUR. COMM’N (Sep. 9, 2025), <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai> (characterizing the GPAI Code of Practice as a “voluntary tool, prepared by independent experts in a multi-stakeholder process, designed to help industry comply with the AI Act’s obligations for providers of general-purpose AI models.”).

314. *Commission Opinion on the Assessment of the General-Purpose AI Code of Practice*, EUR. COMM’N (Aug. 1, 2025), <https://digital-strategy.ec.europa.eu/en/library/commission-opinion-assessment-general-purpose-ai-code-practice>.

315. AI Act, *supra* note 2, art. 55(2) (“Providers of general-purpose AI models with systemic risk may rely on codes of practice within the meaning of Article 56 to demonstrate compliance with the obligations set out in paragraph 1 of this Article, until a harmonised standard is published.”).

316. *Id.* (“Providers of general-purpose AI models with systemic risks who do not adhere to an approved code of practice or do not comply with a European harmonised standard shall demonstrate alternative adequate means of compliance for assessment by the Commission.”).

317. Paul Nemitz & Amin Oueslati, *How US Firms Are Weakening the EU AI Code of Practice*, TECH POL. PRESS (June 30, 2025), <https://www.techpolicy.press/how-us-firms-are-weakening-the-eu-ai-code-of-practice/>.

318. *Id.*

319. Gian Volpicelli & Kurt Wagner, *Meta’s Kaplan Signals Pushback Against EU Regulation for AI*, BLOOMBERG (Feb. 4, 2025), <https://www.bloomberg.com/news/articles/2025-02-04/meta-s-kaplan-signals-pushback-against-eu-regulation-for-ai>.

Commission and the AI Board finalized the process that put the Code in place.³²⁰ But this story shows that companies try to find where the path of least resistance is, institutionally.

E. IN WHICH IT ALL COMES DOWN TO POWER POLITICS

This tracing of power politics brings us to our final point. What happens if U.S. and Chinese AI companies decide to ignore the Act?³²¹ Arguably, there are requirements in the Act with which it is impossible for providers of generative AI to comply. Will companies forego the EU market, or will EU regulators water down or fail to enforce the law?

Another way of framing this question is to ask what role European regulators will play in the development and deployment of technology this time around. European data protection law famously has been exported around the world.³²² But the AI Act is different. First, the AI Act enters the world stage at a very different standing than European data protection law. EU data protection law, as discussed, had substantial historic legitimacy within member states. Moreover, it involved repeat existing players, including regulated companies accustomed to its values, institutions, and mechanisms. Second, data protection law was deliberately designed for export: it contains the (in)famous “adequacy” mechanism, in which Europe will not export EU persons’ data unless a country has been found to have adopted adequate data protection law.³²³ (The United States represents a notorious exception.)

The AI Act, by contrast, was imposed top-down. There is increasing internal EU pressure to deregulate, to be more competitive.³²⁴ The AI Act contains no adequacy mechanism; it relies, instead, on the harmonizing effects of technical standards. Those come, as mentioned, with their own map of power politics, including initial salvos by the United States and increased involvement by China. And the Trump administration, including the erstwhile ally Elon Musk³²⁵ and AI booster JD Vance, have been playing transnational

320. *Commission Opinion*, *supra* note 314.

321. Professor Selbst calls this “the Principle of ‘Fuck You.’”

322. See ANU BRADFORD, *THE BRUSSELS EFFECT* (2020); Anupam Chander, Margot E. Kaminski & William McGeveran, *Catalyzing Privacy Law*, 105 MINN. L. REV. 1733 (2021).

323. Paul M. Schwartz, *Global Data Privacy: The EU Way*, 94 N.Y.U. L. REV. 771 (2019); see also GDPR, *supra* note 19, art. 45 (discussing transfer to third countries).

324. MARIO DRAGHI, REPORT ON THE FUTURE OF EUROPEAN COMPETITIVENESS (2024) (commonly referred to as the “Draghi report”), https://commission.europa.eu/topics/eu-competitiveness/draghi-report_en.

325. Adam Satariano, *E.U. Prepares Major Penalties Against Elon Musk’s X*, N.Y. TIMES (Apr. 3, 2025), <https://www.nytimes.com/2025/04/03/technology/eu-penalties-x-elon-musk.html>.

realpolitik with European regulators.³²⁶ This time around, there may not be a Brussels Effect in which Europe exports its values.³²⁷ Rather, the AI Act risks being gutted through standards from the inside out—or, being gutted in implementation or amended.

V. CONCLUSION

The AI Act already feels like a regulation from another era. It is long, complex, and grounded in a legal regime unfamiliar to a U.S. audience: the NLF. By framing AI systems through product safety law, EU lawmakers aimed to encourage the uptake of AI in the EU. But they ended up creating an instrument—what we've called a legal exoskeleton—that could end up hollowed out through its center, leaving real human rights work to other EU laws and institutions. Time will tell.

326. J.D. Vance, *Remarks by the Vice President at the Artificial Intelligence Action Summit in Paris, France*, AM. PRESIDENCY PROJECT (Feb. 11, 2025), <https://www.presidency.ucsb.edu/documents/remarks-the-vice-president-the-artificial-intelligence-action-summit-paris-france> (“[W]e believe that excessive regulation of the AI sector could kill a transformative industry just as it’s taking off, and we’ll make every effort to encourage pro-growth AI policies [W]e need international regulatory regimes that foster[] the creation of AI technology, rather than strangle[] it. And we need our European friends, in particular, to look to this new frontier with optimism rather than trepidation.” (cleaned up)).

327. See Marco Almada & Anca Radu, *The Brussels Side-Effect: How the AI Act Can Reduce the Global Reach of EU Policy*, 25 GER. L.J. 646 (2024).