U.S. Chamber of Commerce

1615 H Street, NW
Washington, DC 20062-2000
uschamber.com

September 11, 2023

Geoffrey Gertz
Director for International Economics
National Security Council
The White House
Washington, DC 20500

Benjamin Della Rocca
Senior Policy Advisor
National Economic Council
The White House
Washington, DC 20500

Dear Director Gertz and Mr. Della Rocca:

We seek assistance and engagement from the Biden administration regarding artificial intelligence (AI) legislation being considered by the European Union that could deviate from a risk-based approach and established principles of good regulatory practice. If adopted as drafted, the EU's regulatory regime could undermine efforts to establish responsible standards for AI and market interoperability. The U.S. Chamber of Commerce has identified several critical concerns that would significantly impact US scientific and technological interests and undermine our industrial superiority. Your help in addressing the Chamber's concerns outlined below are imperative to protect U.S. scientific and technological interests and to maintain a thriving global ecosystem for innovation.

Please note, this letter and the concerns outlined therein are specific to the EU AI Act, and they do not speak to any other legislation or governance.

Our key concerns include:

- Burdensome targeted requirements for general-purpose AI systems, foundation models, and generative systems irrespective of any concrete risk could impede innovation and hinder the adoption of beneficial technologies, posing a threat to US competitiveness.
- Far-reaching prohibitions and high-risk classifications limit AI's transformative potential and would stifle market growth, directly impacting the ability of US companies to thrive in the European market.
- The imposition of unilateral export restrictions raises concerns about restrictions on international collaboration, hindering US access to global markets.
- The extensive access that regulators may be granted to companies' source code poses significant risks to IP protection and cybersecurity, threatening proprietary technology developed by US companies.
- Limiting specific provisions to very large online platforms or so-called gatekeepers is highly discriminatory and not justified.
- Greater clarity is needed on the implications for companies that use applications based on foundation models.

The trilogue process is already underway between the Council, Parliament, and Commission with the aim to finish by the end of this year. We are concerned that an overly ambitious timeline will not provide the time needed to agree on proportionate solutions.

A unilateral approach by the EU could set precedents for other market economies as we are already seeing in Brazil and Canada. The same goes for non-market economies, which could exploit these rules to their advantage and threaten the future competitiveness of U.S. industry as well as the EU digital single market. Additionally, moving forward in this direction will undermine regulatory collaboration between the EU and its trading partners.

Our concerns are set out in detail below, and we urge U.S. Government interlocutors to raise these concerns with your European counterparts.

Definitions of AI Systems

It is imperative that the EU refine its definition of high-risk AI systems, ensuring a narrower scope that focuses on specific contexts and risks not already covered by existing regulations. An overly broad approach may have significant difficulty differentiating between AI and less complex programs. In this regard, we note that the Parliament has proposed using the OECD definition of AI. The Chamber strongly encourages the trilogue to adopt this approach as the U.S. has already done using the OECD definition of AI for NIST's Risk Management Framework. This will facilitate the EU adopting a definition that is consistent with international standards and avoid creating unnecessary barriers to trade.

The EU should ensure development of technology specific definitions for general-purpose AI, generative AI, and foundation models to provide legal clarity aligned with relevant standards and principles (i.e., OECD). Copyright-specific solutions should be defined separately from the AI Act, if required. The trilogue should also consider existing mechanisms on text and data mining included in the EU Copyright Directive and jurisprudence such as on search engines and indexes. This would avoid overburdening the AI-specific negotiations.

Prohibited AI Systems

The prohibition of AI used for biometric identification (BID) is an area of divergence between the Commission, Council, and Parliament. The Parliament's proposed blanket ban on BID in public spaces to tackle risks of mass surveillance would outlaw beneficial use cases and risks hampering this enabling technology. The scope of such a ban would be disproportionately wide and we urge against this option. Many products provided by private entities use AI for biometric identification, e.g., personalized voice assistants. Many of them allow multiple individuals to use them, e.g., several family members. However, this does not qualify these products as "mass surveillance" tools and they cause no harm, especially considering overlapping requirements from the GDPR that already apply. This includes, for example, features to provide access to people with disabilities, for smart homes, entertainment, and other routine provisions of personalized services.

The Commission proposal and Council text would also ban the use of 'real-time' remote BID for the purpose of law enforcement in all cases except a specific set of purposes such as targeted search for crime victims or prevention of imminent threats to life and safety. However, the

Parliament text would ban all uses of 'real-time' remote BID in publicly accessible spaces by law enforcement authorities as well as private organizations. The scope of this ban would be disproportionately wide, and we urge against this option as well.

The Chamber supports the Commission and Council approach, i.e., acknowledging the risks to fundamental rights posed by government use of AI for surveillance purposes, but also recognizing the vital public safety and national security benefits available from responsible deployment of AI in specific cases paired with strict and meaningful safeguards. These include clearly defined processes and controls such as human review, sufficient confidence scoring, judiciary supervision, clear use policies, reasonable boundaries around data retention, and transparency measures.

Classification of AI Systems as High Risk

The Chamber proposes that the EU keep the Council's proposal to exempt systems with purely accessory function. This is an important classifier to enable the widespread adoption of machine learning systems across the economy. These are machine learning systems that do not make decisions that could significantly impact people's lives. They are often used in conjunction with other systems, such as decision-making systems, but they do not make the final decision themselves. Classifying these systems as high-risk would place an unnecessary burden on U.S. businesses and could discourage the development and use of these beneficial technologies.

It is also important to maintain the European Parliament's key suggestion providing that systems covered under Annex III will be classified as high-risk only if they pose a significant risk of harm to the health and safety. This would strengthen the risk-based approach and limit instances where systems pose only limited risk.

Finally, the EU should reject the notification system as proposed by Parliament as this provision would effectively constitute a pre-marketing authorization procedure for a large number of AI systems. The resulting bureaucracy risks creating burdens and significant delays, hampering beneficial innovation. Further, such a procedure would be inconsistent with product safety legislation more broadly and create strain on EU authorities and discourage U.S. companies from bringing systems to the European market.

Targeted Requirements for General Purpose AI (GPAI)

Imposing targeted requirements on all GPAI systems, including foundation models and generative systems, regardless of risk could have the unintended effect of depriving the EU of access to essential low-risk AI systems that improve people's lives. GPAI systems are an important part of the AI ecosystem and have democratized access to and use of AI technologies for a wide variety of organizations. Designating risk-and purpose-neutral AI technologies as an entirely new class of AI systems would fundamentally affect the architecture of the AI Act, undermining the goal for a carefully balanced, practical, risk-based, and effective approach. Alternatively, reverting to the AI Act's initial approach, which focused on specific use cases of applications with the potential to cause significant, irreversible harms, would allow innovation in low-risk, general-purpose AI technologies to flourish.

GPAI systems, by their nature, do not possess a predefined 'intended purpose,' which means developers may not be aware if a customer plans to deploy the AI in a high-risk scenario or modify it, for example, by training the system with new data, adding new features, or integrating it into another AI system. In these situations, it is important that the right allocation of responsibilities across the AI value chain is carefully considered, and compliance obligations fall on the party best positioned to comply based on value chain considerations. While AI developers may not always be able to anticipate all potential risks, it is important that they provide clear guidelines on the recommended use cases and limitations of the GPAI system. At the same time, AI deployers play a crucial role in assessing and managing risks associated with the intended use of the AI system. There are, however, circumstances where the GPAI is both developed and deployed by the same entity, providing an opportunity for more comprehensive control and risk assessment. In all situations, it is essential that the regulation of GPAI includes tailored and technically feasible requirements, grounded in existing standards and interoperable principles, such as those from the OECD.

Targeted Requirements for Foundation Models (FM)

Targeted requirements for foundation models, as proposed by the European Parliament, would apply whether or not the FM could be used for a high-risk system. The selected requirements also include copyright-specific obligations for FMs that are used as Generative AI, e.g., disclosing training data protected under copyright law. Several of these requirements for FMs impose high compliance burdens. Some are technically not feasible, e.g., training only with specially designed datasets. This would de facto ban the development of FMs that are large language models (LLM) as LLMs are essentially trained on the internet. It is crucial to return to a risk-based approach and that any targeted requirements for foundation models be proportionate, technically feasible, and aligned with international standards and interoperable principles, e.g., those agreed at the OECD or the recent White House Voluntary AI commitments.

Recent discussions on regulating FMs also suggest that the EU's "digital sovereignty" agenda is creeping into the AI Act. The proposal supports the introduction of tailor-made requirements for providers of foundation models (independent of their risk profile), but also recommends a distinction between "smaller" foundation models and "systemic" models, with lower burdens for the former. This asymmetric approach to regulation would be in line with the Digital Services Act (DSA), which imposes more stringent rules on Very Large Online Platforms (VLOPs) and Very Large Online Search Engines. The suggested criteria of distinction could include the amount of money invested in the model, the amount of computer usage, and capabilities. Smaller models would still be subject to strict requirements, but the most stringent obligations would be placed on systemic model providers. This proposal specifically targets American companies building frontier models, and is anti-competitive.

The Parliament's approach focusing certain requirements on the model layer via its FM concept is helpful. However, some of the specific obligations that the Parliament places on foundation models are too far reaching and burdensome. The Parliament's definition of foundation models should be clarified to reflect that such models are "intended to be" adapted and "integrated into a range of different downstream applications." Additionally, to align the regulation at the model level more closely to a risk-based approach, providers of foundation models should be

exempt from requirements when they have explicitly excluded all high-risk uses listed in Annex III, e.g., in the instructions of use or information accompanying the foundation model (borrowing a concept that appears for GPAI models in the Council's Article 4c).

Requirements placed on foundation model providers should take into account relevant existing legislation as well as the technological realities of the AI value chain. Further, they should be calibrated to model level risks, be practicable, and only extend to what foundation model providers can reasonably address during design and development. Thus, rules allocating responsibilities should not only be risk-based, but aligned better to where the decisions are made regarding such risks. Absent changes in the approach, it appears likely that the AI Act will concentrate obligations at the infrastructure/model layer – where American companies are the leading developers – rather than the application layer, where more of the EU user community is implicated. In the final text, it will be crucial to both ensure that high-risk obligations are feasible and that they are allocated appropriately across the value chain.

In their amendments, the Parliament has proposed new rules for providers of FMs and operators who specialize in generative AI systems. Such rules were not included in the Commission's initial text, and thus have not been subject to thorough legal and economic assessment, despite their potential far-reaching impact on the EU AI ecosystem. Some of these are very unclear and could lead to legal uncertainty. For instance, the obligation for FM providers and their users to publish a summary of the copyright-protected training data. It is not clear to whom the summary needs to be disclosed, or which operators the obligation should apply to. Furthermore, this provision risks endangering the trade secrets of providers and rights holders. For this reason, we recommend that the final text of the AI Act does not include such rules.

Value Chain Responsibility

The terminology of the AI Act, e.g., "provider" and "user" as proposed by the Commission and Council does not sufficiently distinguish between roles in the AI value chain (i.e., AI developers, deployers, end users, and other actors), or provide clarity on which parties are responsible. Companies want to understand clearly how and when to comply with the legal requirements. Deployers are usually best suited to know whether their use case will be high-risk, but would struggle with development-related requirements. Developers are typically not well placed to attest use case related requirements.

To ensure legal clarity, obligations should primarily address the actor best placed to comply with requirements. This is in most cases the deployer. Deployers should in return have a legal right to require that developers make contractual commitments to effectively assist them with compliance where required. This would include the developer's acknowledgement of the use of its AI in a high-risk system, the specific obligations developers would be responsible for, and the developer's obligation to assist with requests by a national authority. This mechanism would not apply in situations where the GPAI developer has clarified in the instruction of use that the GPAI must not be used in a high-risk use case (see suggestions on GPAI). This balanced approach inspired by Article 28 of the GDPR would ensure compliance by the deployer enabled by the developer. Additionally, the terminologies should be clarified, particularly between AI developers (i.e., those who make available AI, pre-trained models and

the like), deployers (i.e., those who implement an AI system or who deploy the use case for end-users), and end users (i.e., consumers or those using AI for personal use), and their respective responsibilities. The Council as well as the Parliament both include helpful proposals in this context which should be further specified.

Alignment with Sectoral Legislation for Single Conformity Assessments

Many products with AI components will be governed in parallel by horizontal legislation (AI Act) and by sectoral legislation (e.g., the EU Medical Device Regulation). To avoid double conformity assessments for high-risk products under horizontal and sectoral legislation, the European Parliament introduced Article 8.2a. This Article should be supported as it stipulates for high-risk AI systems that fulfil the requirements of the AI Act under the Union harmonization law listed in Annex II, Section A, their requirements and obligations of the AI Act related to high-risk AI systems shall be deemed fulfilled. The relevant conformity assessment shall also be carried out as part of the procedures laid out under Union law listed in Annex II, Section A. As a result, the AI Act must be aligned with sectoral legislation to avoid unnecessary administrative burdens and costs, such as two separate technical documentations.

List of High-Risk Use Cases in Annex III

The Parliament's proposed considerable expansion of the list of high-risk AI use cases without proper evidence weakens the EU AI Act's risk-based approach. This is especially true for the classification of recommendation systems of VLOPs that are social media and biometric systems.

Recommender systems of VLOPs should not be classified as high-risk. This would be discriminatory against U.S. companies, as recommender systems are used by a large number of websites. Many recommender systems are used to help users find relevant content. The Parliament's proposal also lacks a definition of social media and thus would create legal uncertainty.

The initial proposal of the Commission that classifies all biometric identification systems as high-risk should be rejected. Discrimination on the basis of biometrics is already prohibited under EU law. In fact, biometric systems can be used to prevent discrimination by helping to identify and remove illegal content, such as child sexual abuse material. The vast majority of AI in BID is not deployed in sensitive areas, but enables routine provision of services that make life easier or that entertain.

The Parliament proposes to expand the high-risk category to also include "biometric-based systems" and "AI systems intended to be used to make inferences about personal characteristics". These terms would significantly expand this category and include a broad range of use cases that may indirectly link to biometric data, including non-personal data. Most of these use cases are not relevant for potential discrimination (e.g., BID used for entertainment or in smart home applications). The proposed additions would significantly hamper this enabling technology.

The vast majority of AI in the workplace does not cause harm, on the contrary it enhances safety and drives efficiency. For the most part, AI used in the workplace focuses not on

individuals but on general workplace processes. AI at the workplace may also be used to improve customer experience. “Personal traits” may be used by an AI system to match a customer service agent with a specific customer demand, e.g., considering language skills. Defining all these AI uses as high-risk would impede beneficial use cases. The EU AI Act should exclude use cases that do not track individual traits from the high-risk category of “monitoring and evaluating of performance and behaviour.” Otherwise, U.S. companies would be burdened when applying AI to improve efficiency and safety of work processes.

We recommend that the final text of the EU AI Act mostly maintain the Council's approach to these categories of AI systems. This will ensure a more targeted risk-based approach that does and not unduly restrict the use of beneficial technologies or discriminate against U.S. companies.

While we understand and share policymakers’ concerns around how AI might impact electoral processes in the future, the addition of this use-case to Annex III is unintentionally broad. There are other aspects of the AI Act concerning the labelling of deepfakes, as well as ongoing debates around synthetic media in the context of the EU Code of Practice on Disinformation that will help address concerns around the use of AI in elections. It is also worth pointing out that the Digital Services Act, as the overarching regulatory framework dealing with harmful content, already includes provisions on mitigating negative effects on civic discourse, electoral processes, and public security. These approaches are more appropriate avenues to address policymakers’ concerns and are more aligned with the risk-based approach of the AI Act, compared to a blanket inclusion under Annex III.

Detection of Emotional State of Persons

The detection of emotional state of persons is a use case that has many applications posing low or no risk. The European Parliament introduced a high-risk classification for systems that use biometric or biometric-based data to make inferences about personal characteristics or emotions. This overly broad category would capture a number of low-risk and highly beneficial applications, such as:

- A photo application on a mobile phone that detects the best photo from a photo burst based on whether the subjects have their eyes open or are smiling.
- An application called "guided frame" that helps visually impaired persons take well-framed pictures.
- Systems that can help detect and correct bias.
- Systems are used to detect and reduce bias in other AI systems.

The Council’s position struck the right balance by subjecting systems that use biometric or biometric-based data to make inferences about personal characteristics or emotions only to the transparency obligations in Article 52 of the AI Act. This would allow for the continued development and use of these systems for beneficial purposes, while also ensuring that they are used in a responsible manner.

The European Parliament also introduced in its version a ban on use of Emotion Recognition in the workplace. Neither the Commission, based on its comprehensive risk analysis, nor Member

States proposed a workplace prohibition in their versions of the AI Act. The workplace prohibition will undermine anticipated and already deployed employee health and safety measures. It will also hinder innovation and competitiveness in European manufacturing, retail, logistics, and services industries.

Access to Intellectual Property, including Source Code, Algorithms, and Data Sets

Under the AI Act, European regulators may have the authority to demand access to businesses' data, source code, and algorithms. While there may be precedent for this practice in certain limited circumstances, this is a broad regulatory authority without important safeguards. Requirements on providing access to source code and other proprietary technology or information—if ultimately deemed necessary in exceptional, high-risk use cases—should adhere to EU and international law that protects commercially sensitive information. Regulators having access to a company's privately held datasets and AI systems' source code will expose valuable intellectual property, trade secrets, and personal information to cyberattacks and industrial espionage, including from adversarial countries. Similarly, obligations to retain datasets for a duration longer than required for their intended use need to be thoroughly weighed against data privacy concerns, standards for data minimization, and cybersecurity best practices. This is indicative of a broader trend in Europe that devalues company investments in data and data-driven innovations.

Conclusion

We appreciate the U.S. Government's support in urging the European Union to address these serious concerns to safeguard the interests of U.S. companies, which requires refining its definition of AI systems to narrow its scope and address specific risks not covered by existing regulations; striking a balance between protecting IP and granting regulated access to ensure cybersecurity and preserve the innovative capabilities of US companies; and promoting international collaboration, avoiding unilateral export restrictions, and fostering regulatory harmonization with trading partners.

We welcome the opportunity to speak with you about these issues in greater detail and will contact your respective offices to schedule a meeting. In the meantime, please don't hesitate to reach out with any questions. Thank you for your consideration of our views.

Sincerely,

Marjorie Chorlins
Senior Vice President, Europe
U.S. Chamber of Commerce