



U.S. CHAMBER OF COMMERCE

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

June 12, 2020

Response to the European Commission’s *White Paper on Artificial Intelligence: A European Approach to Excellence & Trust*

The U.S. Chamber of Commerce welcomes the opportunity to provide comments on the European Commission’s *White Paper on Artificial Intelligence: A European Approach to Excellence & Trust* (“White Paper”).

The U.S. Chamber of Commerce (“Chamber”) is the world’s largest business federation, representing the interests of more than three million enterprises of all sizes and sectors. The Chamber is a longtime advocate for stronger commercial ties between the United States and the European Union. According to a recent Chamber study jointly commissioned with AmCham EU, the U.S. and EU are together responsible for over one-third of global gross domestic product, and transatlantic trade and investment supports 16 million jobs on both sides of the Atlantic.¹ The Chamber is also a leading business voice on digital economy policy, including on issues of artificial intelligence, data privacy, cybersecurity, digital trade, and e-commerce. In the U.S. and globally, we advance sound policy frameworks that support economic growth, promote consumer protection, and foster innovation.

The Chamber believes in AI’s potential as a force for good to raise human productivity and expand economic opportunity to benefit consumers, businesses, and all of society. Last year, we issued ten principles for policymakers who are considering action on artificial intelligence:²

1. Recognize Trustworthy AI is a Partnership
2. Be Mindful of Existing Rules and Regulations
3. Adopt Risk-Based Approaches to AI Governance
4. Support Private and Public Investment in AI Research and Development

¹ U.S. Chamber of Commerce & AmChamEU, [The Transatlantic Economy 2020](#).

² U.S. Chamber of Commerce, [Artificial Intelligence Principles](#).

5. Build an AI-Ready Workforce
6. Promote Open and Accessible Government Data
7. Pursue Robust and Flexible Privacy Regimes
8. Advance Intellectual Property Frameworks that Protect and Promote Innovation
9. Commit to Cross-Border Data Flows
10. Abide by International Standards

We hope these principles serve as a reference point for the European Commission (“Commission”) as it considers a future AI governance framework. Comprehensive and substantive consultations with the business community, including a thorough examination of the requirements and enforcement mechanisms proposed, is necessary to ensure that any regulatory intervention does not result in a burdensome or ineffective regime. The Chamber’s comments on the White Paper are below.

Understanding Europe’s Push for “Technological Sovereignty”

Before delving into the specific issues raised by the White Paper, the Chamber would like to address the Commission’s broader push for “technological sovereignty.” Representing companies that are heavily invested in Europe and for whom the EU represents a major market, the Chamber shares the Commission’s goal of advancing the European digital economy, building the digital skills of Europe’s workforce, and preparing Europe’s industrial base for a data-driven future.

At the same time, however, we are concerned about the EU’s drive towards “technological sovereignty.” We welcome a strategy that improves European capacity to compete and attract investment, provided the EU’s approach does not erect barriers or otherwise discriminate against foreign-headquartered companies in the name of creating national champions. **The EU’s efforts to advance the use and development of AI must not shut it off from the rest of the world, as its future competitiveness depends on the ability of all businesses, regardless of size or sector, to remain connected to and engaged with the global economy.** Restrictions on the use of technology developed outside of the EU risks disadvantaging Europe’s own AI capacity, as many businesses benefit from partnerships with non-EU organizations, including those providing cloud capabilities and AI-related components, datasets, and software. **We therefore call on the Commission to explicitly disavow approaches to AI governance that may inhibit market access or disadvantage non-European providers of AI technologies and applications.**

An Ecosystem of Excellence

The Chamber believes in AI's capacity to address many of humanity's most pressing challenges, from climate change to the COVID-19 pandemic, and for delivering much-needed growth at a time of substantial economic uncertainty. According to the McKinsey Global Institute, AI may deliver an additional \$13 trillion of output to the world economy by 2030.³ Given that AI is an evolving and diverse suite of technologies that continues to scale and innovate, a stable regulatory environment and continued investment from governments and the business community is essential. **As the Commission looks to promote Europe's economic recovery and boost its competitiveness, policymakers should prioritize AI investments and stronger incentives for data-driven innovation, rather than focusing solely on establishing an ambitious and overly restrictive regulatory framework for "high risk" AI applications.** We also recommend that the Commission's broader proposals on data governance, including those in the *European Strategy for Data*, incorporate principles of contractual freedom, voluntary data sharing, and non-discriminatory approaches for data transfers, rather than such policies as mandated participation in data pools.

The Chamber is encouraged by several of the proposals outlined in the Commission's "ecosystem of excellence" and applauds the aim of unifying efforts by Member States. We agree that "the private sector should be fully involved in setting the research and innovation agenda and provide the necessary level of co-investment" in the European Union. We recommend that the EU focus in particular on: investing in an AI-ready workforce; helping small and medium sized enterprises understand and use AI; research and development; public sector adoption of AI; and opening government data sets for use by AI developers.

International Aspects

The Chamber strongly supports cooperation between the Commission and the United States Government ("USG") on AI policy to strengthen the transatlantic digital economy and avoid regulatory divergences that inhibit market access and hamper future innovation. We note the important role played by the Commission, the USG, and likeminded nations such as Japan in developing the Organization for Economic Cooperation & Development's ("OECD") *Recommendations on Artificial Intelligence* and their subsequent endorsement by the G20 in the *2019 Ministerial Statement on Trade and the Digital Economy*. With the release of the USG's draft *Guidance for Regulation of Artificial Intelligence Applications*, the U.S. has initiated its own process for articulating a risk-based AI governance framework. **We urge the Commission to engage the USG and the**

³ McKinsey Global Institute, *Notes from the Frontier: Modeling the Impact of AI on the World Economy*, 2018.

U.S. business community to advance the interoperability between these two emerging frameworks. Indeed, cooperation between the EU, the U.S., and likeminded countries such as Japan is necessary to face the common challenge posed by non-market economies that exploit illegal state subsidies, rely on forced technology transfers, and undermine fundamental human rights.

The Chamber welcomes the Commission's commitment to address data localization and other restrictions on data flows in the context of bilateral trade negotiations and at the World Trade Organization ("WTO"). **The ability to move data across borders and access information will determine the speed at which AI technologies can be developed and used in the global economy.** Policies that restrict data flows constitute market access barriers that will diminish AI-related investment and innovation and limit access to AI technologies. With this in mind, we encourage the Commission to embrace high-standard provisions on digital trade, such as those outlined in the *United States-Mexico-Canada Agreement* and the *United States-Japan Digital Trade Agreement*. If replicated in the EU's bilateral agreements and at the WTO's Joint Statement Initiative, these provisions will effectively address protectionist policies such as data localization.

An Ecosystem of Trust

Public trust in AI technologies is necessary to advance their responsible development, deployment, and use. **Given the speed and complexity of technological change the Commission and European Member States cannot foster trust in AI alone; rather, a partnership between governments, the business community, and other stakeholders is needed.** The Chamber recognizes the joint efforts by the Commission and the business community to pilot an assessment list for ethical AI, which was founded on the work of the Commission's High-Level Expert Group on Artificial Intelligence. Multi-stakeholder initiatives such as these are important for advancing voluntary and flexible AI governance frameworks. Indeed, the OECD's AI Policy Observatory has catalogued similar partnerships around the world to address issues of public concern related to artificial intelligence.

The Commission's White Paper asserts that, "...lack of trust is a main factor holding back a broader uptake of AI" and that a "...clear European regulatory framework would build trust among consumers and businesses in AI." The White Paper, however, fails to provide sufficient evidence for these observations. While consumers may not fully comprehend the full potential of AI, consumers interface with businesses using AI at scale every day. This suggests consumers may be more comfortable with AI than even they realize. In defining the "problem" to be addressed

by a future regulation, the White Paper outlines a series of material and immaterial harms that *may* result from the use of AI. **What remains unclear is: 1) the degree to which the Commission has documented the prevalence of harms across a sufficiently representative number of AI use cases; and 2) how the harms contemplated may be appropriately addressed through regulatory action, as opposed to voluntary, multi-stakeholder initiatives already pursued by the Commission and others around the world.**

An evidence-based approach to policymaking is essential. With this in mind, the Chamber encourages the Commission to integrate into its future proposals for AI governance the relevant principles of scientific integrity, information quality, and robust technical analysis. These principles are necessary to verify the degree to which there is a trust deficit in AI among the European public, the nature of the concern, and how this may be addressed in a targeted and balanced manner. Activities performed, and decisions aided by AI are often already accountable under existing European law, including the *General Data Protection Regulation* (“GDPR”) and numerous sectoral requirements. **If the Commission hastily pursues AI regulation without an appropriate base of evidence, it risks stifling Europe’s development and use of this vital technology, undermining the continent’s drive towards greater competitiveness.**⁴

A Definition of Artificial Intelligence

The Commission’s definition of AI will prove critical in determining the scope of its proposed governance framework. While the White Paper does not formally propose a definition, the Commission describes AI as “a collection of technologies that combine data, algorithms and computing power.”⁵ While the Chamber agrees with this description in a general sense, it underscores the difficult task that the Commission faces, as all contemporary software is a collection of data, algorithms, and computing power. An overly broad approach may have significant difficulty in differentiating between AI and a less complex program.⁶ We encourage the Commission to pursue a narrower definition of AI and to focus on the subcategory of AI systems which, in specific contexts, may present risks that are not already adequately addressed by existing regulations. We further encourage the Commission to ensure that such a definition excludes traditional AI programs that operate according to hard coded rules, and instead focuses on future AI systems that are capable of learning on their own.

⁴ Information Technology & Innovation Foundation, *Who is Winning the AI Race?*, 2019.

⁵ We note that work of the High-Level Expert Group on Artificial Intelligence, including [A Definition of AI: Main Capabilities & Scientific Disciplines](#).

⁶ A children’s calculator, for example, qualifies as a technology combining data, algorithms, and computing power.

The European Union's Existing AI Governance Framework

The Commission should undertake a thorough and comprehensive review of EU and Member State laws and regulations that already provide a governance framework for AI across different sectors.⁷ Failure to appropriately account for these rules before instituting a new governance framework may lead to overlapping and contradictory obligations in areas as diverse as financial services, healthcare, transportation, and data protection. In the case of automated vehicle (“AV”) technology, AI regulations would potentially conflict with and duplicate work by the Commission under the revised *General Safety Directive 2019/2144* and the development of AV standards at the United Nations Economic Commission for Europe. Likewise, AVs are fully covered by existing European and Member State liability regimes, offering no basis for future reforms for this AI application. A review of all existing European laws, regulations, and frameworks relating to AI would be consistent with recommendations adopted by the Commission and Member States at the OECD and G20.⁸ Importantly, this exercise would also enable the Commission to focus on areas where existing laws may need to be modified or removed to enable the development, deployment, and use of AI in the single market.

Recommendations for a Risk-Based Approach

As described in our *Principles on Artificial Intelligence*, the Chamber recommends that governments incorporate risk-based approaches rather than prescriptive requirements into frameworks governing the development, deployment, and use of artificial intelligence. The Chamber agrees with the Commission that there can be “no one-size-fits-all rules that can properly accommodate the many unique characteristics of every industry making use of this technology and its impact on individuals.” Indeed, we believe that AI use cases that involve a high risk should face a higher degree of scrutiny than a use case where the risk of concrete harm to individuals is low. The Chamber recommends that the risk-based approach outlined in the White Paper be improved in the following respects:

First, the Commission should consider how a future governance framework would account for the variety of risks that may or may not be applicable across different AI contexts. A risk-based approach considers factors such as safety and human life, impact on critical infrastructure, financial market stability, and the capability to cause concrete harm to individuals. **A binary approach that relies, even in part, on categorizing entire sectors as “high-risk” or “not high-risk” is insufficiently nuanced and unable to capture the heterogeneous nature of “risk.”** Moreover, the

⁷ We acknowledge the important [contributions](#) of the Expert Group on Liability and New Technologies in this emerging debate, which we address in our comments on liability below.

⁸ OECD AI Recommendations, 2.3(b); G20 Trade & Digital Economy Statement, *ibid*.

definition of industry sectors is not clear-cut, as these sectors constantly evolve, and AI applications may often be used across different sectors.

Second, the Commission’s risk-based approach should recognize the different factors that are important for understanding the risks that AI may pose in a specific context. As written, the White Paper focuses on the *severity* of potential harms. **The Chamber recommends that any future risk analysis also be proportional and based on additional factors, including the *probability* and *scale* of harm.** The same AI application, put to the same use, will pose different risks depending on the way it is integrated into businesses’ operations, as the degree of human oversight and safeguards such as monitoring may vary between contexts and may increase or decrease the risk based on the technology in question. The risk factors around AI are also a function of mitigations available, whether technical, application, or process based. The Chamber therefore recommends that a future framework also recognize risk mitigation.

Third, the framework should explicitly recognize tradeoffs, such as the opportunity cost of *not* adopting AI and whether the benefit of using AI in a specific context outweighs its harm. **Any risk assessment should account for the significant social, safety, and economic benefits that may accrue when an AI application replaces a human action.** The point of comparison for any risk assessment of an AI application should be the solutions that are currently available, rather than an imagined standard of perfection. AI systems will have issues that can be mitigated and controlled, but existing processes and solutions, including those based on human actions, have their own issues and challenges. The aim of AI technologies is to improve on the status quo, not to be perfect.

Fourth, the White Paper contemplates “exceptional uses” of AI which would always be treated as “high risk” regardless of sector or context, including recruitment, labor issues, remote biometric identification, and consumer protection. These are open-ended categories and the *a priori* designation undermines the entirety of the Commission’s risk-based proposal, which should focus instead on specific use cases. **We strongly caution the Commission against including these provisions in its future proposal.** An expansive designation of “high risk” uses may capture a significant number of AI applications, thereby imposing potentially burdensome obligations across the single market. Such an approach would undermine, rather than support, the EU’s economic competitiveness.

Fifth, the proposal would incorporate so-called “immaterial damages” into its consideration of risks and AI harms. We recommend that this provision be qualified, as it has a potentially broad scope. Any regulation of AI should be specific, narrowly

tailored to appropriate use cases, and weighed against the economic and social benefits forfeited by its enactment.

Proposed Requirements for “High-Risk” AI

The Commission outlines six possible obligations for companies whose AI applications qualify as “high risk,” which may be applied *ex ante* or *ex post*. In addition to comments on the specific proposed requirements, the Chamber offers the following overarching points regarding obligations on “high risk” applications of AI.

As noted above, **the Commission should avoid prescribing a single set of requirements, because this approach fails to account for distinctions between the types of risk presented by different applications.** While some may present a high-level of risk to individuals’ privacy, others may present a safety risk, meaning that the potential degree of harm is different. The Chamber recommends the Commission establish the requirements as guidelines whose implementation will necessarily vary from case-to-case. Such an approach would enable companies that develop and deploy AI across the technologies’ lifecycle to implement risk management practices in a way that best fit the use case and risk profile.

Second, the Commission’s proposal should recognize the importance of non-regulatory approaches to governing “high-risk” AI. **Non-regulatory approaches often achieve the same policy objectives and offer the same level of protections as regulatory approaches, but without many of the burdens and unintended consequences.** For example, the development of voluntary consensus standards on the national and international level is a highly effective means of addressing the challenges and opportunities presented by emerging technologies such as artificial intelligence.⁹ Similarly, multi-stakeholder initiatives have the greatest capacity to identify gaps in AI outcomes and to mobilize AI actors to address them, including through the development of tools such as algorithmic impact assessments. We further encourage the use of voluntary codes of conduct as accountability tools in the AI lifecycle, and not only for AI applications that do not meet the criteria of “high risk.”

I. Requirements on Training Data

Fairness and non-discrimination principles are essential for establishing public trust in AI. In assessing the impact of an AI application on fairness and non-discrimination, including the role played by training data, the Chamber recommends that the Commission consider several factors. First, the principles of fairness and non-

⁹ Cf: OECD AI Recommendations, 2.5(c); G20 Statement.

discrimination are not unique to AI, so it is important that the Commission consider how the principles are applied in existing human-based contexts, including in existing sectoral and Member State laws. New approaches to these principles should not supersede established definitions and practices, but instead should focus on identifying harms that could potentially arise and be empirically linked to discrimination. Second, it is important to note that eliminating or removing bias from models may not be technically possible in all circumstances.

The Chamber supports an approach that considers mitigation as a tool to address and reduce bias. However, **it is unclear how legal requirements on training data can be constructed to produce fair or non-discriminatory outcomes.** The proposed obligations for developers to “ensure datasets are sufficiently representative” is impractical. It is unclear how to determine what is “sufficient” — especially for providers of multipurpose AI systems—as there is no clear or widely accepted definition or metrics for datasets. Such a requirement may conflict with GDPR, under which developers are not meant to have access to sensitive attributes like ethnicity. We also note that it is possible to create a high performing model even using biased, low quality training data, and the reverse is also true. Rather than putting requirements on training data, it would be better to focus on AI outcomes and examine whether they are within an acceptable range, since this will ultimately determine the real world impact of an AI system.

II. Requirements to Keep Records and Data

The Commission proposes record-keeping obligations for any use of AI that is classified as “high risk.” As with other requirements, this should be formulated in a flexible manner to allow for a wide variety of contexts and delivery formats. Strict record-keeping may be more appropriate for some use cases rather than others. Additionally, retention of datasets may conflict with obligations under GDPR to minimize and/or delete personal data. Such an obligation may even undercut the privacy benefits of on-device processing because it would effectively force companies to centrally collect and store personal data used in training algorithms. It would also prevent the use of off-the-shelf, open-source models, since developers will generally have no access to the data used to train them. Finally, the requirement may present challenges for copyrighted datasets authorized for only short-term access.

III. Requirements to Disclose Information

The Chamber believes that disclosure about the nature and function of AI is appropriate in high-risk contexts. It is important, however, that a disclosure requirement is not taken to an extreme. General information regarding the kind of

datasets used to train an algorithm, such as its age or size, or what checks it has been subjected to (e.g., for bias) may prove useful. At the same time, in business-to-business contexts, information provision should be limited to parties with a legitimate interest. Any requirement put forward by the Commission must respect intellectual property rights and avoid mandatory disclosure of detailed data or information which reveals AI algorithms or the underlying code, as this may violate business confidentiality and undermine an AI's safety and reliability by opening it up to attacks by adversarial parties.

IV. Requirements to be “Robust and Accurate”

The goal of incentivizing risk assessments to “ensu[re] that AI systems are robust and accurate” is shared by many stakeholders. However, it is unclear how an abstract standard such as this could be mandated through a legal requirement. **At present, there are no clear or widely accepted conceptions about what constitutes a “robust and accurate” AI system, nor are there mature standards against which systems might be measured.** If a stringent requirement, such as full traceability of every outcome of an AI system, was mandated, it would in practice restrict AI systems to an extremely limited, basic set of techniques (e.g., static decision trees). Mandating specific techniques legislatively may inadvertently undermine longer term safety by discouraging organizations from developing improved approaches. In fact, “accuracy” may not be the end goal of every AI system. In the financial services sector, for example, businesses may reduce the accuracy of an AI system purposely to increase false positives, thereby flagging more transactions as potentially risky. In doing so, businesses are taking extra steps in conducting due diligence to combat fraud. Any legal requirement demanding absolute accuracy would run counter to these practices. These concerns may be avoided if the Commission focuses more on promoting favorable outcomes and less on the inner workings of AI technology. Separately, the Chamber requests greater clarification on the scope of the Commission’s approach to “reproducibility.”

V. Requirements for Human Oversight

The Chamber agrees in a general sense that human oversight of AI applications in high-risk contexts is good practice. However, **different degrees and approaches of oversight are needed across different AI use cases.** Human oversight when training an AI system is fundamentally different from human oversight in AI used in a real-time operation. Requiring an AI system’s output to be reviewed by a person before being acted upon may make sense for some applications (e.g., AI systems used for critical, non-time-sensitive medical diagnostics). For other applications, though, it could undermine their safety (e.g., requiring human intervention for driverless vehicles), lead to sluggish output, reduced privacy, or undermine accuracy (e.g., if human reviewers

lacked the necessary expertise or were more biased). At an extreme, it could even put people at risk by delaying automated safety overrides. Requirements for human intervention must therefore be clear, narrow, and adapted to the realities of different sectors. The Commission should also consider how requirements for human intervention may impede businesses' ability to provide innovative goods and services to consumers and to compete with businesses in jurisdictions where it is not mandated.

VI. Requirements for Remote Biometric Identification

The Chamber welcomes the Commission's commitment to launch a broad debate on the uses of remote biometric identification. We believe that these tools, such as facial recognition technology, have enormous potential to enhance security and safety, and enable innovation across a wide variety of sectors including transportation, retail, hospitality, and financial services. Already, the technology can be used in applications including airline passenger facilitation, criminal investigations, theft prevention, and fraud detection. As such, we caution policymakers against pursuing overly burdensome regulatory regimes, such as moratoriums or blanket prohibitions. We note that uses of remote biometric identification are already regulated under the *General Data Protection Regulation*. We recommend that any requirement for remote biometric identification tools recognize these existing requirements and remain consistent with the Commission's future AI governance framework, that is, a risk-based approach with specific requirements triggered only in those use cases designated as "high risk."

Recommendations for a Compliance & Enforcement Regime

The Commission proposes a conformity assessment requirement for AI applications designated as "high risk." The assessment may include procedures for testing, inspection, or certification to verify that a series of mandatory requirements applicable to high-risk applications are met. The Chamber strongly advises against instituting a separate conformity assessment regime for AI beyond that already required for existing technologies into which AI is incorporated. We also submit the following points:

First, **a new conformity assessment regime would likely serve as a significant bottleneck on the development and deployment of AI in the EU**, as companies would need to win approval from regulators before deploying AI-enabled goods and services in the Single Market. Many innovative small and medium sized enterprises that may have neither the time nor resources to undergo such a process will either avoid investing in perceived "high risk" areas or deploy their solutions abroad. The additional costs will reduce competition and choice in the Single Market for AI goods and services deemed as "high risk."

Second, a **conformity assessment regime raises significant trade and intellectual property concerns, as companies will be reticent to allow an outside organization or government agency to inspect an algorithm and datasets used in its development.** Algorithms and datasets are often proprietary, and their protection is enshrined in the EU’s existing trade obligations. The *Japan-European Union Economic Partnership Agreement* explicitly prevents the transfer of source code as a condition for market access.¹⁰ **The Chamber cautions against any requirement that “high risk” AI goods and services be retrained in the EU as a condition for market access.** There is no guarantee that EU datasets or training that takes place on European soil will do anything to improve the performance of an AI system. Blocking the use of foundational non-European datasets would risk reducing a system’s performance and could even exacerbate the risk of discrimination. Any such requirement would violate the EU’s WTO national treatment obligations and, possibly, its commitments under the *General Agreement on Trade in Services*.

Third, a **combination of *ex-ante* risk self-assessment by companies using “high risk” AI and *ex-post* enforcement would likely achieve similar outcomes to a conformity assessment requirement within much faster timeframes.** Such an approach is already enshrined in GDPR, which requires data protection impact assessments for potentially high-risk use cases, with Data Protection Authorities (“DPAs”) providing backstop enforcement. Finally, If the Commission proceeds with a conformity assessment proposal, it should also take into account whether existing products on the market would have to retroactively undergo inspection, as this may potentially create a significant backlog for newly established testing centers.

A Voluntary Labeling Scheme

Voluntary labeling may be a promising approach to promote trust in AI systems. As written, though, the Commission’s proposal to develop such a scheme for AI applications not designated as “high risk” leaves a number of important questions unanswered, including: the authority that would assess and issue the label, the enforcement mechanism, and whether it would create a *de facto* conformity assessment requirement for all AI systems. The Chamber believes that more work needs to be done in a multi-stakeholder setting to arrive at broad agreement on the standards that would be covered by a voluntary labeling framework. We recommend that a framework not be binding on developers or deployers who opt to use the label, as this would disincentivize participation by many businesses.

¹⁰ Japan-European Union Economic Partnership Agreement, Article 8.73.

Recommendations for a Governance Framework

In contemplating an AI governance framework, the Commission should integrate existing regulatory structures as much as possible. Regulators—whether in financial services, pharmaceuticals, transportation, or safety and security—are best placed to interpret and apply a risk framework to a specific context. At the same time, for the purposes of legal certainty, DPAs should retain their existing competencies over AI applications processing personal data.

The insurance industry exemplifies one among many sectors that are already governed by a robust framework of laws, regulations, and standards. Insurers are closely supervised by regulators to ensure proper consumer protections and risk management techniques are in place. As a result, they apply the same level of care and diligence to AI as they have done with other third-party products and services that have been integrated into their operations over decades. The Commission should be mindful of imposing new requirements for AI onto this already robust regulatory and supervisory framework. Hampering the ability of insurers and other sectors to leverage AI at scale may come with significant opportunity costs.

The Commission must also avoid the creation of a patchwork of AI regulations across the European Union. We encourage the Commission to account for the lessons learned from GDPR’s implementation. As noted in our recent submission on the regulation’s two-year review process, GDPR’s implementation has been plagued by inconsistent interpretation by DPAs, weak coordination at the European level, and persistent attempts to undermine the one-stop-shop function. Consequently, GDPR *in practice* has served to fragment the EU’s data protection landscape, rather than unify it as originally intended. We are concerned that these issues may be repeated and exacerbated under a potential AI governance framework, as it will necessarily involve more regulators balancing more legitimate public policy aims, including data protection. To ensure the smooth functioning of a future AI governance framework, it is critical for businesses to interact with one regulatory authority on issues of enforcement and compliance.

The Chamber welcomes the Commission’s statement that “the governance structure should guarantee maximum stakeholder participation,” including in the implementation and development of a future AI governance framework. **The Commission should provide ample opportunities for the business community to participate in all stages of the policymaking and rulemaking process, including when designating specific sectors and / or applications as “high risk.”** This is necessary for strengthening accountability and improving regulatory outcomes, and for reducing adverse consequences on the ability of businesses to develop, use, and deploy AI systems in the European Union. In addition, we seek to further strengthen the

business community's partnership with European institutions to educate them on the perception, opportunities, and impacts of AI applications.

The European Union's Liability Regime

The White Paper includes an important and necessary focus on potential liability issues arising from the development and use of artificial intelligence. The Commission considers that certain adjustments should be made to EU and national liability regimes to adapt them to AI development, potentially including adjustments to the usual burden of proof, and expansion of strict liability to new situations.

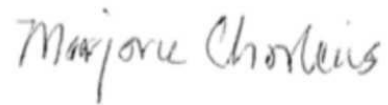
Such adjustments may seem like natural precautions when much remains unclear about how AI will develop. **However, the Commission should be mindful that fault should continue to be the guiding principle underpinning liability. Imposing impossible burdens or presumptive liability in relation to future unknown risks could greatly undermine the incentive to innovate.** Specifically, where the intrinsic purpose of AI is autonomous decision-making, there is a limit to what organizations can realistically foresee, plan for, and mitigate against. If current doctrines are crudely adapted to require developers to account for unforeseen or unforeseeable circumstances, this may reduce their appetite to allow AI autonomy, which essentially voids its ultimate purpose.

An imbalanced liability regime would therefore risk penalizing organizations for matters genuinely outside their control, or even deter them from developing societally useful AI in the first place. **Overall, the Commission must be careful not to intervene too hastily to adjust liability principles in nascent markets, as it is difficult to anticipate their evolution.** Given the novelty of the issues involved and the legal uncertainty surrounding AI, developing soft measures such as codes of conduct and guidance, in close collaboration with relevant stakeholders, is more appropriate than the introduction of new or any immediate changes to existing hard measures.

Conclusion

The Chamber thanks the Commission for the opportunity to provide these comments. The U.S. business community is engaged in significant trade and investment with the European Union and is proud of its continued contributions to our vibrant bilateral commercial relationship. We look forward to continued dialogue on the Commission's AI proposal, as well as other foundational digital policy issues.

Sincerely,



Marjorie Chorlins
Senior Vice President
European Affairs
U.S. Chamber of Commerce



Sean Heather
Senior Vice President
International Regulatory Affairs
U.S. Chamber of Commerce