THE

EU AI Act and Brussels Effect.

How will American AI firms respond to General Purpose AI requirements?

Claudia Wilson

August 2024

CENTER FOR AI POLICY

Executive Summary

When there's no international regulation on a sector, who sets the global rules of the road? Sometimes, domestic regulation remains domestic and companies simply adhere to local laws. However, there is a trend of American companies complying with EU regulations *outside* of the EU. Policy researchers have labeled this phenomenon the Brussels Effect.

This paper asks whether the requirements of the EU AI Act, which has just been enacted, will lead to a Brussels Effect for American businesses. The EU AI Act is a comprehensive attempt to mitigate the safety and ethical risks associated with AI. Its requirements are significantly more involved than any US legislation, which predominantly consists of voluntary commitments. Thus, any Brussels Effect would likely elevate the safety and ethical protocols within American AI companies. We focus specifically on the requirements for General Purpose AI (GPAI) models, which are more recent additions to the Act, and thus have not been featured in as much research on the Brussels Effect. In Exhibits 2 and 3, we summarize each GPAI requirement, the cost implications of those requirements, and the likely response of GPAI companies.

The Center for AI Policy (CAIP) finds that large American companies are likely to remain in the EU market and be generally compliant with the EU AI Act, even when operating in the US. The potential revenue in the EU will, in most cases, exceed the costs of compliance, causing most firms to remain in the EU market, though some smaller firms may choose to leave. As we explain in Section 4, market forces may already such as the cost of compute may be causing smaller companies to struggle in GPAI markets.

Most firms will choose to be compliant with the EU AI Act, because it is likely to be more profitable than running two separate compliance processes and potentially training two separate models. Many of the EU AI Act requirements implicate practices prior to or during training such that if companies wished to avoid EU regulation outside the EU, they would need to incur the expensive cost of training twice. Several of the requirements, such as documenting the training process or reporting capabilities, are associated with fixed costs *across* markets, meaning that there is no additional cost to being compliant in the US as well as the EU.

Firms may choose not to adhere to incident tracking and cybersecurity standards outside the EU since non-adherence could potentially offer some minor savings. However, companies are likely to already conduct incident tracking and have cybersecurity protocols, so any cost benefit of avoiding EU levels of regulation is limited. Additionally, bifurcating processes across countries would introduce complexity, making it unlikely that firms would take up this opportunity to avoid EU regulation.

Despite widespread debate about the "opt-out" copyright requirement, we find that, with some EU facilitation, this is feasible and firms will likely choose to adhere to this requirement outside of the EU. If the EU clarifies what the best practice is for declaring an "opt-out", such as through a central list like Spawning AI's Do Not Train Registry or a consistent robots tag, firms will have greater clarity about how to adhere to the requirement. While this "opt-out" requirement may limit or increase the cost of training data, it seems from agreements with media companies that

US AI companies are already preparing to spend more on copyright training data, regardless of the EU AI Act. Given that adherence occurs prior to training, firms would need to train two separate models to avoid this requirement in the US and stay in the EU market. Thus, the "opt-out" requirement is likely to have a Brussels Effect.

If companies were to have separate products for separate markets, it would be due to the evaluations requirement delaying launch dates. However, this is a tenuous incentive. Evaluations have a fixed cost across markets, so would have no additional cost to being globally compliant. There would be additional costs associated with running two separate evaluations processes, as companies would need to maintain two separate models. Thus, since cost doesn't incentivize separate evaluations, companies would need to care greatly about timing to be non-compliant with EU evaluations requirements.

These findings may be particularly relevant to those interested in US leadership on global AI standards. While the US can continue to influence the *de jure* international AI agreements, it takes time to achieve international consensus. If the US remains at its current political impasse regarding AI policy, its companies will be *de facto* regulated by the EU. To be a leader in global AI governance, the US should consider moving ahead with concrete domestic AI safety policy.

This paper is structured as follows: Section 1 provides context on the EU and US AI regulatory environments. Section 2 defines the Brussels Effect and provides other examples of it in the digital regulation space. Section 3 outlines the framework companies would likely use to choose their response to the EU AI Act. Section 4 analyzes the incentives for American companies to remain or exit the EU GPAI market. Section 5 examines how each of the Act's GPAI requirements will incentivize companies to pursue "differentiation", where they don't adhere to the Act outside of the EU, or "non-differentiation", where they do adhere to the Act in the US.

Contents

1. Context on EU and US regulatory environments

The EU and the US are known to provide vastly different regulatory environments, and the AI policy space is no exception. Typically, the EU has more stringent regulation with a focus on civil liberties and protections, while the US favors an open business environment that incentivizes innovation. As of 2024, there is a clear divergence between the US and EU AI regulatory environments.

The EU AI Act (henceforth referred to as "the Act") is a comprehensive attempt to ensure that "AI developed and used in the EU is trustworthy, with safeguards to protect people's fundamental rights"[1]. It has a series of requirements based on risk tiers for more specific-use AI tools and was amended to include specific requirements for General Purpose AI (GPAI) models following the release of large language models (LLMs) such as ChatGPT. After years of deliberation, the Act was enacted August 2024. Some measures, such as prohibitions of unacceptable risk AI will come into effect after six months. Enforcement of General Purpose AI (GPAI) requirements will not commence until August 2025 for new entrants and August 2027 for existing market participants.

Conversely, the US has predominantly focused on introducing voluntary commitments for AI companies at the federal level[2]. Recent bills, such as the Artificial Intelligence Research, Innovation, and Accountability Act (AIRIA), propose self-certification for AI companies, but have not yet specified what exactly would be required and whether companies would be independently audited[3]. Although States such as California and Colorado have begun to consider more binding AI legislation, the federal landscape remains largely unregulated[4].

2. Historic examples of the Brussels Effect and digital regulation

Anu Bradford defines the Brussels Effect as the penetration of European rules and regulations into many aspects of economic life outside of Europe[5]. This "unilateral regulatory globalization" is distinct from multilateral regulation, where multiple countries input and agree upon rules of the road[6]. The Brussels Effect can occur through a *de facto* effect, where companies choose to adhere to EU regulation outside of the EU, or through a *de jure* effect, where other countries implement EU-inspired regulation.

Bradford identifies examples of the Brussels Effect in antitrust, chemical regulation, environmental protection, and food safety[7]. In the case of antitrust, there is Brussels Effect due the non-divisibility of mergers. In other words, international companies find it difficult to merge in one country, but not another. Thus, the EU successfully blocked the Honeywell and General Electric merger. Conversely, the scale economies of chemical production enabled the Brussels Effect of the Registration, Evaluation, Authorisation and Restriction of Chemicals (REACH) Act. Several companies including Ikea, Lego, Mattel, Revlon, Unilever and L'Oreal changed their global production process to be consistently compliant with REACH. Similarly, farmers have found it difficult to fully stratify their supply chains into genetically modified organisms (GMO) and non-GMO products, so many have decided to become globally adherent with European

GMO regulation. In the case of environmental protection, some international firms adhere to the 2003 Restriction of Hazardous Substances Directive (RoHS) because it's simpler to have a single compliance process.

There are also several examples of the Brussels Effect occurring recently with digital regulation. As outlined below, there has been consistent adherence to the General Data Protection Regulation (GDPR) and Digital Markets Act (DMA), although both regulations have been met with some resistance by specific companies. This resistance takes the form of non-compliance in the EU or releasing different versions of products in different markets.

The GDPR is one example of a "de facto" Brussels Effect occurring in the digital regulation space, albeit with patchy compliance and some resistance by Big Tech. Europe introduced the GDPR in 2018 to harmonize and strengthen data privacy laws across the EU[8]. Preferring to maintain access to the EU's 450 million potential users, many firms chose to adhere to the GDPR requirement inside and outside of the EU, resulting in cookies consent requests globally[9]. This widespread compliance to a GDPR requirement indicates some degree of a Brussels Effect.

While the GDPR has had a clear Brussels Effect, it has been met with resistance by some companies and has not seen full adherence. While Meta gave EU citizens the chance to "opt-out" of their social media accounts being used as training data, US citizens were given no such option[10]. In July 2024, Meta also announced that it would not release multimodal AI models in the EU due to lack of clarity around GDPR requirements[11]. This choice follows a back-and-forth between Meta and the EU. Meta claims to have briefed the EU on its plan to train models on publicly available social media posts, allowing users to opt-out[12]. After receiving minimal feedback and then publicly announcing these plans, Meta was ordered to pause training in June and answer more questions regarding data privacy[13]. Prior to its decision to limit EU products, Meta had also been fined €1.2 billion for sending European citizens' data to the US[14]. Meta's response seems, in part, a protest against GDPR and, in part, driven by a desire to meet launch dates elsewhere. These decisions to bifurcate "opt out" decisions and limit features in the EU represent a limit to the Brussels Effect.

Similarly, the DMA has seen a mixed Brussels Effect. The 2024 regulation seeks to address anti-competitive behavior and the dominance of Big Tech[15]. Some companies, such as Whatsapp, have chosen to become globally compliant. The messaging app, which previously did not enable chatting with other apps, has changed its interoperability, though is waiting on other apps to respond[16]. Conversely, Apple has been accused of violating the DMA and is limiting features available in the EU, actively resisting the Brussels Effect. Although Apple claims to have made changes to comply with the DMA, it is under investigation by the European Commission for not allowing app developers to promote alternative distribution channels on the App Store[17]. Apple has also chosen to delay the EU rollout of generative AI features because of DMA interoperability requirements, which the European Commission Vice-President labeled a "stunning declaration" of anti-competitive behavior[18].

Finally, the Digital Services Act (DSA) is yet to see a major Brussels Effect. Also a 2024 introduction, the DSA "attempts to impose transparency in areas like algorithms and advertising, fight online harassment and disinformation, protect minors, stop user profiling, and eliminate dark patterns (design features intended to manipulate our choices on the web)"[19]. Several researchers have commented that the DSA's requirements may directly conflict with American interpretations of free speech[20]. Academics predict that if there is a Brussels Effect it would be modest in scale[21].

In summary, the GDPR, DMA, and DSA all have some degree of a *de facto* Brussels Effect, but none of the regulations has seen full compliance outside the EU, and some firms have even been accused of violating requirements inside the EU.

3. Framework for US companies' decision making

3a. Remaining in the EU market

For any AI company operating in the EU market, there are two key decisions. These are articulated in Exhibit 1 below and are based on Anu Bradford's and the Centre for Governance of AI's research[22]. Our research extends their work by focusing explicitly on the GPAI requirements of the Act, which were introduced after Bradford and the Centre for Governance of AI conducted their research on AI regulation and the Brussels Effect.

The first decision companies make is whether to remain in the EU market or to exit. This framework assumes that companies remain in the EU market as long as their EU operations are profitable, that is, if their EU revenue exceeds their EU costs, including cost of compliance and expected spend on fines.

3b. Differentiation vs non-differentiation

If companies choose to remain in the market, they then must choose whether to have a single EU-compliant product that spans multiple markets ("non-differentiation") or whether to bifurcate their products and only adhere to EU regulation in the EU market ("differentiation"). For example, one form of differentiation would be for a company to train one model for the EU and one model for the US, each with their own dataset. Alternatively, a company could train a single model, but run different evaluations or different incident tracking for the EU or the US. Thus, there are multiple degrees of differentiation that companies can pursue, each at different stages in product development.

The choice to differentiate is determined by which option generates greater profits outside the EU. If it is more profitable to sell an EU-compliant product outside the EU, firms will do so. If it is less profitable to sell an EU-compliant product outside the EU, firms will train and maintain two separate models. Companies are more likely to choose non-differentiation if the market is attractive, demand for products is inelastic, and the cost for differentiation is high.

Exhibit 1: Differentiation vs non-differentiation profits [23]

	Non–differentiation profits outside the EU		Differentiation profits outside the EU
Revenues	Revenue from selling EU compliant products outside the EU		Revenue from selling non-EU compliant products outside the EU
	—	≥	—
Costs	Variable compliance cost of producing an EU-compliant product for non-EU markets		Additional regulatory costs of non-EU compliant products (fixed and variable costs) Duplication costs of two products

4. Why large companies will remain in the EU

The EU AI Act is unlikely to completely erode the profitability of the EU market for large GPAI companies, thus they will choose to remain in the EU market. Since the EU has not released impact assessments for the GPAI requirements as it did for other portions of the Act, we do not have estimated costs of compliance for these specific requirements[24]. However, we can use the impact assessment of the "high risk" AI systems as a rough indicator of whether the Act will bankrupt companies - the main reason why companies would exit the market. These impact assessments estimated that companies which don't already fulfill *any* requirements would spend an additional 17% of revenue on compliance. We expect that most, if not all, GPAI companies are at least partially fulfilling these requirements and thus 17% would be an overestimate of cost. Even so, comparing an overstated cost burden of 17% to AI profit margins of 50-60% still yields a profitable, albeit less profitable, company[25]. Again 17% is likely an overestimation, but regardless, AI companies are unlikely to leave due to profitability.

Some smaller AI companies may choose to leave the EU. After the GDPR, Verve, a mobile marketer, and online games companies Gravity Interactive and Uber Entertainment exited the market due to compliance costs[26]. The EU AI Act will undoubtedly cost more than companies' existing compliance processes, particularly for those that are not already conducting some of the required measures. Smaller companies may not have the scale or the profitability of an Anthropic, which supposedly has a gross profit margin of 50-55%[27], and thus they may not be able to bear the burden of compliance.

While consolidation of frontier AI companies would be a bad outcome, regulation would not be the sole cause. The cost of AI components has quickly built a barrier to entry for cutting-edge GPAI models. Computing power is wildly expensive with GPT-3 taking $50-$100 million to train[28]. Talent is also not cheap. Many large tech companies also have unparalleled access to data. Even as algorithms become more efficient and computational performance improves, driving costs down, larger companies will simply be able to outspend

smaller start-ups and dominate the cutting edge of GPAI[29]. Thus, there may be exits and consolidation of smaller AI companies, but this will be caused by market factors rather than the EU AI Act.

While AI companies will factor in the risk of violation, we expect that companies will not choose to remain in the EU with the intent of violating the Act. The Act has high penalties for violating GPAI requirements. Companies face fines up to 3% of global revenue or EUR 15,000,000, whichever is highest. Since these fines scale with firm size, we expect that large AI companies will try their best to adhere to the Act if they do remain in the market. Thus, AI companies will estimate a relatively low chance of violation when factoring the cost of penalties into their decision to remain in the EU.

There is a chance that companies exit the EU market to protest regulation, but this is unlikely given the potential revenue forfeited. Many AI companies have expressed their disapproval of the Act and Sam Altman even threatened to withdraw OpenAI from the EU[30]. However, exiting a market is a commitment and cuts off access to future growth. If companies wish to return to a market, it can be expensive to restart operations. Ultimately, the EU is an attractive market, with forecast growth of 33% per annum to 2030 and an estimated 22% of the global AI market[31]. While such estimates can be imperfect, the EU's 31.5 million businesses and working population of 286 million represent a significant market that will likely be worth complying with the EU AI Act[32].

5. Why companies will likely pursue non-differentiation

For those companies that choose to remain in the EU market, it is likely that a Brussels Effect will occur. As we outline below, the EU is an attractive market and non-EU demand is probably only slightly elastic, thus the Brussels Effect is dependent on the cost of differentiation, which we find to be higher than non-differentiation.

The EU is a relatively attractive market for GPAI market size of $6.4b in 2023 and estimated market of $110.8b in 2030[33]. While growth rates, particularly those related to AI, are inherently tricky to predict, the EU's large population and mature industries mean that there is enough demand for companies to be genuinely interested in maintaining presence.

Since it is difficult to back-out elasticity from consumer choices, we assume that American consumers are at most slightly 'elastic' when it comes to GPAI. In other words, it requires a reasonably large reduction in capabilities to encourage substitution away from EU compliant-models. In this situation, elasticity largely depends on customer preferences and what they value in a GPAI model. If American consumers do not value EU compliance and instead are focused purely on price and capability, then they may be more price sensitive or 'elastic'. However, if customers value EU compliance, perhaps seeing the models as secure and private, they may be more willing to absorb higher prices for EU-compliant models. Elasticity matters because it is an indicator of the extent to which non-EU demand will fall if firms pursue non-differentiation and consequently pass on higher costs or limit available features to non-EU customers.

Many of the Act's requirements imply a high-cost of differentiation, because they are fixed costs across markets and / or occur early in the AI development process. Given the high cost of training GPAI models - estimated to be $78 million for GPT-4 and $191 million for Gemini Ultra - companies will only pursue differentiation if there is a clear upside in revenue or large reduction in costs from having non-EU compliant models in the US[34]. Thus, companies are less likely to differentiate if requirements occur early in the model development process, meaning that companies would need to train two models, or requirements have a fixed cost across markets, meaning that there is no additional cost to be compliant with that measure in two markets. There is also another driver of cost for differentiation. If companies were to maintain separate models for separate markets, they would also face a small but ongoing "identification" cost to ensure that their customers are in the correct market. Exhibit 2 and 3 below outline each individual requirement, the nature of the costs associated, and the implications for differentiation preferences.

For example, Article 53a, which relates to documenting the model training process, occurs during training and has a fixed cost across markets, so incentivizes non-differentiation. Similarly, Article 53b requires firms to record and share capabilities, so has a fixed cost across markets. This requirement occurs later in the product development process, so wouldn't necessarily require duplicating training processes to differentiate, but the relatively low fixed cost associated with it means that non-differentiation is likely. Article 53d, which requires a fair effort documentation of the data sources used for training, occurs early in the training process, and would require training two separate models if companies chose non-adherence outside the EU. Some commenters previously aired concerns about the technical feasibility of a highly detailed reflection of data, but subsequent recital 107 of the Act indicates that the template required will be 'simple, effective, and allow the provider to provide the required summary in narrative form'. Thus, we interpret Article 53d as a feasible expectation of companies and thus an incentive for non-differentiation.

It would be relatively easy to differentiate on incident tracking and cybersecurity standards, but it is unclear that doing so would enable significant cost savings. Articles 55c and 55d have variable costs across markets and occur after the expensive training process, so it would be reasonably affordable to differentiate. Yet, firms will likely conduct some degree of incident tracking and cybersecurity measures in the US regardless of EU regulation, so there is limited cost benefit of differentiation. Firms would also weigh this cost benefit against the complexity of stratifying their incident tracking and cybersecurity processes. Firms are most likely to resist the reporting of incidents required in 55c, but due to matters of brand and self-preservation, rather than due to large additional costs. For the cost reasons listed above, firms are likely to follow internally consistent incident tracking and cybersecurity measures.

Exhibit 2: Requirements for GPAI with no systemic risk[35]

Article (simplified)		Commentary	Implication
53a.	Draw up and keep up-to-date the technical documentation of the model, including its training and testing process and the results of its evaluation	• Fixed cost across markets • Early forking since documentation needs to begin during training	Easier to run **single model**
53b.	Draw up, keep up-to-date and make available information and documentation to providers who intend to integrate the GPAI into their AI system. • Model capabilities and limitations. • Information on data used for training, testing and validation, including the type and provenance of data and curation methodologies.	• Fixed cost across markets • Early forking for data information • Depending on level of detail about data, may be a high effort or cost requirement • Later forking for model capabilities and limitations	Easier to run **single model**
53c.	Comply with Union copyright, specifically enable rights-holders to *opt-out* ex ante from their work being used to train models.	• Fixed cost across markets • Early forking since it relates to data collection • Incremental cost may be minimal given that companies are already pursuing deals with rights holders	Incentive for t**wo models** if adherence will drastically reduce training data Depends on feasibility and interpretation of best-efforts
53d.	Make publicly available a summary of training data, according to a template. It should be generally comprehensive in its scope instead of technically detailed[36] • Listing the main datasets used to train the model • Providing a narrative explanation about other data sources used	• Fixed across markets • Early forking since it relates to tracking data collection	Easier to run a **single model**

Exhibit 3: Requirements for GPAI with systemic risk[37]

Article (simplified)		Commentary	Implication
55a.	Perform model evaluation in accordance with standardised protocols, including conducting and documenting adversarial testing of the model with a view to identifying and mitigating systemic risks	• Fixed cost across markets • Later forking since it occurs after model training • Could potentially delay deployment in non-EU markets	Incentive for **two models** if non-EU launch date is important
55b.	Assess and mitigate possible systemic risks at Union level that may stem from the development, the placing on the market, or the use of GPAI models with systemic risk	• Fixed cost across markets • Later forking	Easier to run a **single model**
55c.	Keep track of, document, and report to relevant information about serious incidents and possible corrective measures to address them	• Variable cost across markets • Later forking	Slight incentive for **two models**
55d.	Ensure an adequate level of cybersecurity protection for the general-purpose AI model with systemic risk and the physical infrastructure of the model	• Variable cost across markets with a fixed cost component	Slight incentive for **two models**

Copyright requirements are one potential incentive for differentiation, depending on interpretation and EU facilitation. Article 53c describes a DSA provision for rights-holders to opt-out of their content being used for model training. Usually, copyright laws are determined by the geography of content creation, so in the case where models are trained, but in this case, the Act has specified that any model available in the EU must adhere to EU copyright law regardless of where the model is trained[38]. Language such as "expressly reserved" indicates that the opt-out is an ex-ante provision, so rightsholders can't retrospectively ask for their content to be removed from a model[39]. The ex-ante nature of the provision limits the burden such that AI companies are not expected to retrain models because of this requirement.

An 'opt-out' may be logistically challenging for firms, but this challenge could be resolved if the EU implements a technical standard or clarifies "best practice". The Act has not yet offered clarity on the mechanisms for rights holders to opt-out, simply stating that they should be "machine-readable"[40]. This ambiguity raises questions of how rights holders can reasonably announce that they're opting out. Does a unilateral statement on an author's personal website count? Does the opt-out need to be mentioned in the text itself, for example in

the opening pages of a novel? Some organizations, such as Sacem, a French music society, and Pictoright, a Dutch picture collecting society, have provided draft "opt-out" statements for their members.

However, it is unclear how the EU envisions rights holders technically conveying their "opt-outs". One organization, Spawning, has created a Do Not Train registry that enables machine readable opt-outs and has partnered with Stability AI and Hugging Face[41]. They have also released code that functions as a machine readable opt-out for website owners, but this doesn't extend to individual works[42]. Ideally, the European Commission would specify that it recommends methods like these for "opt-outs" so that companies have clarity around what best practice compliance looks like. Given that "opt-outs" occur early in the AI development process, differentiation would come with the expense of training two models. Thus, if companies have clarity around logistics, they are more likely to pursue non-differentiation.

It is also unclear whether the EU will punish "best efforts" measures that result in accidental inclusions of "opt-out" data. Clearly a safe harbor for companies who use their "best-efforts" would be preferable for companies, while a harsher interpretation may scare companies into leaving the market or using older versions of models. If any accidental slip-up can be slapped with a fine of 3% of global revenue, companies may be unnerved by an impossibly high requirement and potentially ruinous penalties. Given clarifications on other requirements, such as Recital 60k, which explained that documentation of training data can be expressed in narrative form, we think it is likely that the EU will take a reasonable "best efforts" interpretation. Additionally, if the EU introduces a technical standard or clarifies "best practice", then "best efforts" measures are less likely to have accidental inclusions and companies will have greater confidence that they are in compliance with the Act.

The "opt-out" requirement also threatens companies' cheap access to training data, which could interfere with their competitive advantage. So far, cutting-edge innovation by the large AI labs has relied on wholesale access to training data, regardless of whether it is legal or ethical[43]. Thus, it is important to understand, if this opt-out provision is logistically possible, how much data it will affect. This is difficult, because AI companies do not like to disclose their training data, perhaps for the very issue this copyright requirement seeks to address. However, we can make generalizations across datasets. Epoch AI assessed three high-quality datasets and found that they are *"composed of 50% scraped user-generated content (Pile-CC, OpenWebText2, social media conversations, filtered webpages, MassiveWeb, C4), 15-20% books, 10-20% scientific papers, <10% code and <10% news"*[44]. Therefore, if 100% of rights holders exercise their right to opt-out, AI companies may lose 35-50% of their datasets. Of course, it is unlikely that there will be a 100% uptake by rights holders. Even if only 25% of rights holders exercise the right to "opt-out", this means that companies can no longer use 9-13% of their datasets. The magnitude of this effect is augmented by the fact that books and newspapers tend to be the highest quality source of data.

However, the US may also move towards stricter copyright protections and companies are already entering into expensive data agreements meaning that the "opt-out" requirement may not fundamentally change their access to training data. When deciding whether to pursue differentiation, companies will compare the EU copyright requirements to US copyright requirements. There is only an incentive to duplicate models if the US offers a significantly more lax copyright environment. So far, the US has yet to announce any federal legislation to regulate AI companies' use of copyrighted material, but lawsuits are currently underway and will determine whether AI training is "fair use"[45]. In the meantime, large companies are entering into formal agreements with publishing houses, newspapers, and other rights holders. For example, OpenAI's deals with The Atlantic, Vox Media, and at least ten other media firms could be worth over $100 million per year[46]. Google recently signed a deal with Reddit that was "about $60 million per year"[47]. Given how expensive these deals are, AI labs would likely prefer not to have to pay for training data. Yet, the fact that labs have entered these deals voluntarily reveals their expectations about how current litigation may conclude and signal that the era of bootlegging copyright work to train models may be over. Thus, perhaps the EU copyright requirements will only make an incremental difference to the cost of training data compared to the alternative baseline in the US.

The most compelling incentive for differentiation is model evaluation, which could delay US launch dates, but this is a tenuous incentive. Companies are currently competing to be at the cutting edge of innovation and to be recognized by the public for these accomplishments. This means that they are particularly sensitive to launch dates. One could argue that additional evaluations, as required in Article 55a, may delay launch dates and frustrate companies to the point of differentiation. Given that evaluations will have a fixed cost across markets, there is no clear cost incentive to run two separate evaluations processes. It may even be more expensive for companies to differentiate their evaluations, since they would experience the cost of maintaining two models instead of one. Costs such as post-deployment cleanups to correct for bugs and flaws would be duplicated. Thus, the only incentive to differentiate evaluations is related to timing of launch dates.

Depending on a company's existing processes, evaluations could delay launch dates by months. Evaluations could refer to benchmarking or red-teaming or a combination of both. The Act does specify that it will require some form of red-teaming, which typically takes months for complex models such as GPAI[48]. While the Act doesn't explicitly state that it will require benchmarks, they are industry standard, so will likely be required to some extent. One potential required benchmark could be Holistic Evaluation of Language Models (HELM), which assesses metrics such as accuracy, calibration, robustness, fairness, bias, toxicity and efficiency[49]. It can take a day to conduct HELM benchmarking and upwards of $10,000[50]. To be conservative, we assume that companies will be required to conduct multiple of these benchmarks, such that they spend several weeks on benchmarking, though this can be in parallel with red-teaming. OpenAI's CTO claims that they spent "6 months making GPT-4 safer and more aligned", which suggests that any delays to launch caused by EU AI Act evaluations would already be *partially* accounted for in BAU evaluation timelines[51].

AI companies may be able to speed up launch dates by running parallel processes. The Act evaluations will almost certainly require more computing power and resources than the current state, but some of these evaluations can be conducted in parallel, shortening the delay to launch. This delay could be several weeks or several months, depending on the company's internal processes and the specifics of the Act. Ultimately, if the delay is longer and companies believe that it will significantly affect revenue or brand recognition, they will be more willing to bear the cost of maintaining two separate models.

Conclusion

Each requirement of the EU AI Act has different costs associated and thus can imply different incentives. When assessed as a whole, the Act is likely to see some degree of a Brussels Effect, where American companies are complying with the Act even in the US. This adherence is most likely for measures that occur prior to or during training, such as copyright requirements as well as documentation of data sources and the training process. Requirements with fixed costs across markets are also likely to lead to a Brussels Effect, particularly when they have limited cost or revenue benefits for non-adherence in the US, such as documenting capabilities. There are reasonable arguments why incident tracking and cybersecurity measures may incentivize differentiation. These arguments mostly relate to the measures having variable costs across markets and occurring later in the product development process, or potentially delaying launch dates. However, given that US companies already conduct each of these activities to some extent, it is unclear that the EU requirements will lead to a significant increase in cost or time that would outweigh the complexity of running two separate compliance processes or models.

If there is an incentive to differentiate, it would be model evaluations due to potential delays in launch dates. It would be more expensive to maintain two models, so companies would need to highly value timing to pursue this. Companies could also be EU-compliant and accelerate launch dates to some extent by running parallel processes. Their current evaluations timeline will also be crucial in determining how much of a delay the EU AI Act would cause and whether it's worth maintaining two models.

The Center for AI Policy recommends that the EU clarify and facilitate best practices for the "opt-out" requirement. Given the high fines, companies may have concerns about accidental violations of the Act if they don't fully understand the logistical expectations. A technical standard, such as using a robots tag to communicate "opt-outs" or a central "opt-out" registry may clarify to companies what steps they can take to be compliant. There are many reasonable positions that a government could take when deciding how to balance the needs of intellectual property owners against the needs of AI developers, but leaving the EU's policy ambiguous makes all parties worse off.

Finally, the predicted Brussels Effect has interesting implications for US leadership on global AI governance. President Biden's October 2023 Executive Order outlined an aspiration for the US to "lead the way" on "building and promoting (AI) safeguards with the rest of the world"[52]. US Department of Commerce Secretary Raimondo has expressed a desire for the US to be "at the front of the pack" on AI governance[53]. In early 2024, Senators Warner and Blackburn

introduced legislation "aimed at restoring the US's position as a lead in international standards-setting for emerging technologies"[54]. Thus, there is a demonstrated interest for the US to influence international AI governance.

The US is currently pursuing such leadership through participation in multilateral agreements, such as the Bletchley Declaration and the G7 International Code of Conduct[55]. However, these statements are not legally binding and, thus have limited scope to ensure adherence to these commitments. It can be challenging to achieve binding multilateral agreements due to the required consensus of key stakeholders.

US political division and the influence of Big Tech in policymaking has made it difficult to achieve consensus on AI policy[56]. If the US remains at such a political impasse, the fact is, the EU will have greater influence on US AI companies than the US government or any multilateral agreement. As members of Congress consider US AI policy, they should consider whether they want the US to lead on AI governance or be led.

Works cited

[1] European Commission. EU AI Act: first regulation on artificial intelligence. *European Commission*. https://ec.europa.eu/commission/presscorner/detail/en/ip_24_4123 (2024).
[2] White House. Fact sheet: Biden-Harris Administration secures voluntary commitments from leading artificial intelligence companies to manage the risks posed by AI. *whitehouse.gov*. https://www.whitehouse.gov/briefing-room/statements-releases/2023/07/21/fact-sheet-biden-harris-administration-secures-voluntary-commitments-from-leading-artificial-intelligence-companies-to-manage-the-risks-posed-by-ai/ (2023).
[3] Congress. S.3312 - Artificial Intelligence Research, Innovation, and Accountability Act of 2023. *congress.gov*. https://www.congress.gov/bill/118th-congress/senate-bill/3312/text (2023).
[4] Mayer Brown. Colorado governor signs comprehensive AI bill. *mayerbrown.com*. https://www.mayerbrown.com/en/insights/publications/2024/06/colorado-governor-signs-comprehensive-ai-bill (2024); Turner, J., & Turner Lee, N. Can California fill the federal void on frontier AI regulation? *The Brookings Institution*. https://www.brookings.edu/articles/can-california-fill-the-federal-void-on-frontier-ai-regulation/ (2024).
[5] Bradford, A. The Brussels effect. *Northwestern University Law Review*, 107(1), 1-67. *Proquest*. https://www.proquest.com/docview/1286682884/fulltextPDF/2F4AE630ADF347AAPQ/1?accountid=15172&sourcetype=Scholarly%20Journals (2012).
[6] Ibid
[7] Ibid
[8] Burgess, M. What is GDPR? The EU's data protection law explained. *Wired*. https://www.wired.com/story/what-is-gdpr-uk-eu-legislation-compliance-summary-fines-2018/ (2018).
[9] Cookiebot. GDPR cookies, consent, and compliance. *Cookiebot*. https://www.cookiebot.com/en/gdpr-cookies/ (2024).; Lawler, R. Apple may delay AI features in the EU because of its big tech law. *The Verge*. https://www.theverge.com/2024/6/21/24183251/apple-eu-delay-ai-screen-mirroring-shareplay-dma (2024).
[10] Mauran, C. Meta is using your posts to train AI. It's not easy to opt out. *mashable.com*. https://mashable.com/article/meta-using-posts-train-ai-opt-out (2024).
[11] Fried, I. Tech giants up ante by withholding products from EU. *Axios*. https://www.axios.com/2024/07/18/tech-giants-eu-regulation-withholding-products (2024).
[12] Fried, I. Scoop: Meta won't offer future multimodal AI models in EU. *Axios*. https://www.axios.com/2024/07/17/meta-future-multimodal-ai-models-eu (2024).
[13] Ibid
[14] Burgess, M. Meta's $1.3 billion fine is a strike against surveillance capitalism. *Wired*. https://www.wired.com/story/meta-gdpr-fine-ireland/ (2023).
[15] Zorloni, L. The EU is taking on Big Tech. It may be outmatched. *Wired*. https://www.wired.com/story/european-commission-big-tech-regulation-outlook/ (2024).
[16] Ibid
[17] European Commission. Commission sends preliminary findings to Apple and opens additional non-compliance investigation against Apple under the Digital Markets Act. *European Commission*. https://ec.europa.eu/commission/presscorner/detail/en/ip_24_3433 (2024).
[18] Lawler, R. Apple may delay AI features in the EU because of its big tech law. *The Verge*. https://www.theverge.com/2024/6/21/24183251/apple-eu-delay-ai-screen-mirroring-shareplay-dma (2024).; Gkritsi, E. EU Competition Commissioner says Apple's decision to pull AI from EU shows anticompetitive behavior. *Euractiv*. https://www.euractiv.com/section/digital/news/eu-competition-commissioner-says-apples-decision-to-pull-ai-from-eu-shows-anticompetitive-behavior/ (2024).
[19] Zorloni, L. The EU is taking on Big Tech. It may be outmatched. *Wired*. https://www.wired.com/story/european-commission-big-tech-regulation-outlook/ (2024).
[20] Huddleston, J. The Brussels effect?: Potential domestic impacts of international online speech regulation. *Cato Institute*. https://www.cato.org/blog/brussels-effect-potential-impact-speech-regulation-around-world-americans-online-0 (2023).

[21] Husovec, M., & Urban, J. Will the DSA have the Brussels Effect? *Verfassungsblog*. https://verfassungsblog.de/will-the-dsa-have-the-brussels-effect/ (2024).
[22] Siegmann, C., & Anderljung, M. The Brussels effect and artificial intelligence. *Centre for Governance of AI*. https://www.governance.ai/research-paper/brussels-effect-ai (2022). ; Bradford, A. The Brussels effect. *Northwestern University Law Review*, 107(1), 1-67. *Proquest*. https://www.proquest.com/docview/1286682884/fulltextPDF/2F4AE630ADF347AAPQ/1?accountid=15172&sourcetype=Scholarly%20Journals (2012).
[23] Siegmann, C., & Anderljung, M. The Brussels effect and artificial intelligence. *Centre for Governance of AI*. https://www.governance.ai/research-paper/brussels-effect-ai (2022).
[24] European Commission. Study to support an impact assessment of regulatory requirements for Artificial Intelligence in Europe. *European Commission*. https://op.europa.eu/en/publication-detail/-/publication/55538b70-a638-11eb-9585-01aa75ed71a1 (2021).
[25] Gnanaselvam, P. The economics of artificial intelligence (AI) companies. *Medium*. https://medium.com/@praveentgs/the-economics-of-artificial-intelligence-ai-companies-638c496f3745 (2022).
[26] Kottasová, I. These companies are getting killed by GDPR. *CNN Business*. https://money.cnn.com/2018/05/11/technology/gdpr-tech-companies-losers/index.html#:~:text=Europe's%20new%20data%20protection%20law,businesses%20or%20shut%20down%20entirely. (2018).
[27] Gustafson, C. Does AI have a gross margin problem? *Mostly Metrics*. https://www.mostlymetrics.com/p/does-ai-have-a-gross-margin-problem (2024).
[28] Vipra, J., & Myers West, S. Computational power & AI. *AI Now Institute*. https://ainowinstitute.org/publication/policy/compute-and-ai#h-how-is-the-demand-for-compute-shaping-ai-development (2023).
[29] Epoch AI. Key trends and figures in machine learning. *Epoch AI*. https://epochai.org/trends (2023).
[30] Wodecki, B. OpenAI threatens to leave EU over AI Act, then recants. *AI Business*. https://aibusiness.com/nlp/openai-threatens-to-leave-eu-over-ai-act-then-recants (2024).
[31] European Commission. Study to support an impact assessment of regulatory requirements for Artificial Intelligence in Europe. *European Commission*. https://op.europa.eu/en/publication-detail/-/publication/55538b70-a638-11eb-9585-01aa75ed71a1 (2021).; Grand View Research. Europe artificial intelligence market trends. *Grand View Research*. https://www.grandviewresearch.com/industry-analysis/europe-artificial-intelligence-market-report#:~:text=Europe%20Artificial%20Intelligence%20Market%20Trends,33.2%25%20from%202025%20to%202030. (2024).
[32] Eurostat. Business demography statistics. *European Commission*. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Business_demography_statistics (2024).World Bank. Population ages 15-64, total. *World Bank Data*. https://data.worldbank.org/indicator/SP.POP.1564.TO?locations=EU (2024).
[33] Statista. Generative AI in Europe: Market size and outlook. *Statista*. https://www.statista.com/outlook/tmo/artificial-intelligence/generative-ai/europe#market-size (2024).
[34] Maslej, N., Fattorini, L., Perrault, R., Parli, V., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., Wald, R., & Clark, J. The AI Index 2024 Annual Report. *AI Index Steering Committee, Institute for Human-Centered AI, Stanford University*. https://aiindex.stanford.edu/wp-content/uploads/2024/04/HAI_AI-Index-Report-2024.pdf (2024).
[35] EU Artificial Intelligence Act. Article 53: Obligations for providers of general-purpose AI models. *EU Artificial Intelligence Act*. https://artificialintelligenceact.eu/article/53/ (2024).
[36] EU AI Act. Recital 107. *EU AI Act*. https://www.euaiact.com/recital/107 (2024).
[37] EU Artificial Intelligence Act. Article 55: Obligations for providers of general-purpose AI models with systemic risk. *EU Artificial Intelligence Act*. https://artificialintelligenceact.eu/article/55/ (2024).
[38] Riede, L., Talhoff, O., & Hofer, M. The AI Act: Calling for global compliance with EU copyright? *Lexology*. https://www.lexology.com/library/detail.aspx?g=830531fc-c40a-4b26-9dfd-955b37631525 (2024).
[39] Ziaja, G. M. The text and data mining opt-out in Article 4(3) CDSMD: Adequate veto right for rightholders or a suffocating blanket for European artificial intelligence innovations? *Journal of*

Intellectual Property Law & Practice, 19(5), 453–459. https://academic.oup.com/jiplp/article/19/5/453/7614898 (2024).
[40] Keller, P. Protecting creatives or impeding progress? Machine learning and the EU copyright framework. *Kluwer Copyright Blog*. https://copyrightblog.kluweriplaw.com/2023/02/20/protecting-
[41] Spawning. About Spawning. *Spawning AI*. https://spawning.ai/about#content (2024).
[42] Keller, P., & Warso, Z. Defining best practices for opting out of ML training. *Open Future*. https://openfuture.eu/wp-content/uploads/2023/09/Best-_practices_for_optout_ML_training.pdf (2023).
[43] Metz, C., Kang, C., Frenkel, S., Thompson, S. A., & Grant, N. How tech giants cut corners to harvest data for A.I. *The New York Times*. https://www.nytimes.com/2024/04/06/technology/tech-giants-harvest-data-artificial-intelligence.html (2024).
[44] Villalobos, P., Sevilla, J., Heim, L., Besiroglu, T., Hobbhahn, M., & Ho, A. Will we run out of data? An analysis of the limits of scaling datasets in machine learning. *arXiv preprint*. https://arxiv.org/pdf/2211.04325v1 (2022).
[45] Keller, P. Protecting creatives or impeding progress? Machine learning and the EU copyright framework. *Kluwer Copyright Blog*. https://copyrightblog.kluweriplaw.com/2023/02/20/protecting-creatives-or-impeding-progress-machine-learning-and-the-eu-copyright-framework/ (2023). ; Zirpoli, C. T. Generative artificial intelligence and copyright law. *Congressional Research Service*. https://crsreports.congress.gov/product/pdf/LSB/LSB10922 (2023).
[46] AI Policy Weekly. AI Policy Weekly #26: OpenAI inks twelfth major data deal, illustrating steep costs of AI development. *AI Policy Weekly*. https://aipolicyus.substack.com/p/ai-policy-weekly-26?open=false#%C2%A7openai-inks-twelfth-major-data-deal-illustrating-steep-costs-of-ai-development (2024).
[47] Ibid
[48] Sama. Red teaming solutions for generative and large language models. *Sama*. https://www.sama.com/red-teaming-generative-ai#:~:text=The%20duration%20of%20a%20red%20teaming%20engagement,complexity%20of%20the%20model%2C%20the%20desired%20scope (2024).
[49] Guan, X. Navigating the LLM benchmark boom: A comprehensive catalogue. *Holistic AI Blog*. https://www.holisticai.com/blog/navigating-llm-benchmark (2024).
[50] Martineau, K. Tiny benchmarks for large language models. *IBM Research Blog*. https://research.ibm.com/blog/efficient-llm-benchmarking (2024).
[51] Bastian, M. OpenAI's GPT-4 is a safer and more useful ChatGPT that understands images. *The Decoder*. https://the-decoder.com/open-ai-gpt-4-announcement/ (2024).
[52] Biden, J. Executive order on the safe, secure, and trustworthy development and use of artificial intelligence. *The White House*. https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/ (2023).
[53] U.S. Department of Commerce. Biden-Harris Administration announces first-ever consortium dedicated to advancing AI technology and research. *U.S. Department of Commerce*. https://www.commerce.gov/news/press-releases/2024/02/biden-harris-administration-announces-first-ever-consortium-dedicated (2024).
[54] Warner, M., & Blackburn, M. Warner, Blackburn introduce legislation to reestablish U.S. leadership in international standards-setting for emerging tech. *Senator Mark Warner*. https://www.warner.senate.gov/public/index.cfm/pressreleases?ID=EA36C577-070F-412D-9290-3EEF3C961B18 (2024).
[55] The White House. G7 Leaders' Statement on the Hiroshima AI Process. *The White House*. https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/g7-leaders-statement-on-the-hiroshima-ai-process/ (2023); United Kingdom Government. The Bletchley Declaration by countries attending the AI Safety Summit, 1-2 November 2023. *UK Government*. https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration/the-bletchley-declaration-by-countries-attending-the-ai-safety-summit-1-2-november-2023 (2023).;
[56] Lima-Strong, C., & Dou, E. The AI election is here. Regulators can't decide whose problem it is. *The Washington Post*. https://www.washingtonpost.com/technology/2024/06/06/ai-election-2024-us-misinformation-regulation/ (2024).