TECHNICAL REPORT | JUNE 2026

Delays to Frontier AI in the EU and UK

Analysis of 375 LLM Releases by Meta, Google, OpenAI, and Anthropic

John Lidiard, Oleksandra Vereschak, Tom Gibbs, and Markus Anderljung

GovAI

Executive Summary

- In recent years, frontier AI companies have sometimes delayed the release of new models to the EU and UK – or not released them at all. To what extent has this happened, and to what extent were regulatory barriers the cause?
- To investigate this, we collated and analyzed a dataset of 375 LLM releases by Meta, Google, OpenAI, and Anthropic over an eight-year period (June 2018 – May 2026), and used public sources to assess the most likely reason for delays or non-releases.
- By doing so, we hope to inform the design and implementation of AI regulation – in particular, how decision makers weigh regulatory objectives against the risk of models being delayed or withheld from their markets.

How often are models delayed or not released on the EU and UK markets?

- Relative to the US, we find that 11% of model releases were delayed or not released to the EU, and 7% were delayed or not released to the UK.
- The UK saw no delays in accessing the model ranked highest on the Epoch Capabilities Index at any point. The EU faced a 71-day delay for the web-app release of Claude 3 Opus while it ranked as the highest-performing model; API access, which developers use to build applications, was simultaneous with the US.
- Meta had the highest overall rate of delays and non-releases of any company in our dataset, with ~26% of its releases delayed or withheld in the EU and ~15% in the UK.
- We generally see a reduction in delays and non-releases over time. This could reflect companies transitioning from research to commercial organizations, with mature compliance functions and greater incentive to deploy broadly. It may also be due to greater regulatory certainty as EU and UK regulators clarified how data protection regulation applies to LLM-based services.
- In the first five months of 2026 there were no instances of a publicly released model being delayed or withheld from the EU and UK, but there were several instances of models being released to select entities under trusted-access programs (GPT 5.5 Cyber, Claude Mythos, GPT Rosalind). US export controls resulted in Anthropic pulling both the widely available Fable and the trusted-access version of Mythos. In the current environment, US national security concerns – rather than regulation – may pose the greater barrier to European access to frontier AI.

What are the causes of delays and non-releases?

- Regulatory barriers appear to be the primary cause of most delays or non-releases. Of the 68 examples of delays and non-releases in the dataset, we tentatively attribute regulatory factors as the primary, though not necessarily sole, cause of 56.
- Data protection regulation appears to be the main regulatory barrier, especially in relation to training on personal data. Non-text modalities (image, audio, real-time video) seem to face greater barriers than text.
- Although the UK and EU share broadly similar data protection laws, regulatory barriers appear greater in the EU for the period covered by our dataset. This is likely due to the EU's more aggressive regulatory enforcement and slower clarification of how data protection rules apply to the training and deployment of LLMs.
- We find no strong evidence of other regulations, including the EU AI Act, causing delays or non-releases. However, the AI Act's general-purpose AI provisions only came into force in August 2025, covering ten months of our dataset, and will only become enforceable in August 2026 (not covered at all by our dataset).
- In nine instances, the primary cause of a delay or non-release appears to be due to non-regulatory factors, including compute constraints, language support limitations, and broader product readiness concerns.

This work represents the views of its authors, rather than the views of the organization, and does not constitute legal advice. GovAI technical reports have received extensive feedback, but have not gone through formal peer review.

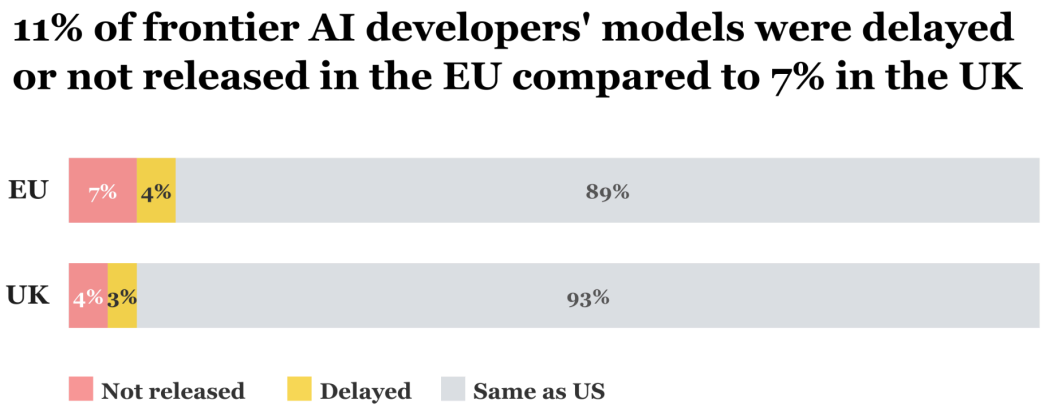

Figure 1: We analyzed a dataset of 375 LLM releases from Meta, Google, OpenAI, and Anthropic (06/2018 – 05/2026) to compare US release dates with EU and UK release dates. We show here the overall percentage of LLMs that were

not released, delayed, or released simultaneously within the EU and UK, relative to the date of their US release. All numbers are rounded to the nearest whole number.

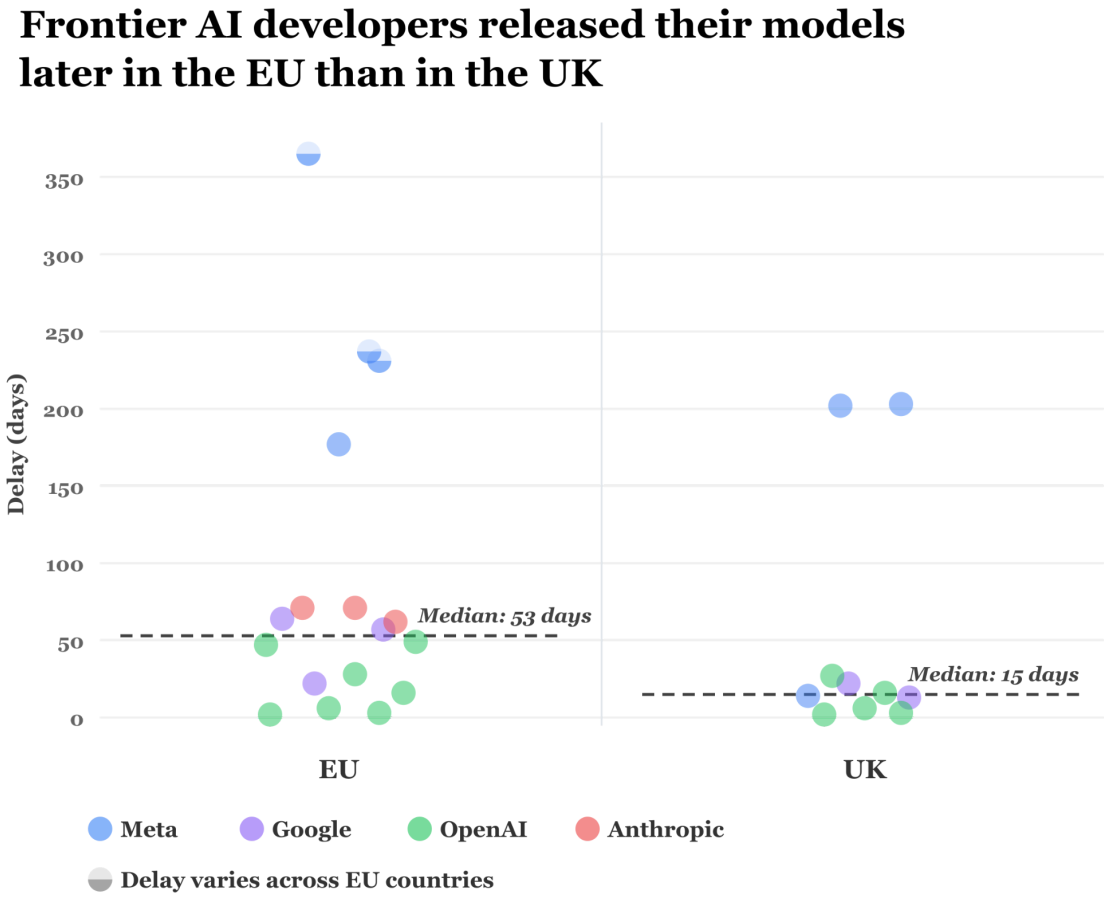

Figure 2: We analyzed a dataset of 375 LLM releases from Meta, Google, OpenAI, and Anthropic (06/2018 – 05/2026) to compare US release dates with EU and UK release dates. Here we show all the instances of a model being delayed to the EU and UK, and median delay per region. Each dot represents an LLM release, which is color-coded according to the company. We use a half-dot where there is variation between EU countries in the release date (two instances).

Delays to Frontier AI in the EU and UK

Analysis of 375 LLM releases by Meta, Google, OpenAI, and Anthropic

John Lidiard[1] *, Oleksandra Vereschak[2]*, Tom Gibbs[3], Markus Anderljung[1] **

[1] GovAI
[2] Future Impact Group Fellowship
[3] Independent Researcher

* - Equal contribution. Authors are free to list themselves as first author on their CVs. For questions and to access the dataset, write to contact@governance.ai

** - Senior author

1. Introduction **4**
2. Delays and Non-Releases of LLMs in the EU and the UK **6**
3. Reasons for Delays and Non-Releases **9**
Meta 19
Google 21
OpenAI 22
Anthropic 22
Methodology **24**
Limitations and Future Work **26**
About the Authors **27**
Acknowledgements **27**
About GovAI **28**

1. Introduction

Overly stringent regulation is often said to delay – or altogether block – access to frontier AI models in Europe.[1] To what extent has this actually played out in recent years?

To answer this question, we compile and analyze a dataset of 375 LLM releases from four leading AI companies – Meta, Google, OpenAI, and Anthropic – assessing how often releases were delayed or withheld from the EU and UK, and whether regulation was the primary cause.

[1] See, for example: Draghi, Mario, "The Future of European Competitiveness," European Commission, 9 September 2024 (criticising regulatory complexity and GDPR/AI Act overlap as a drag on European AI development); "EU needs AI: An open letter on EU regulation," published as an advertisement in the *Financial Times* and coordinated by Meta with signatories including Spotify, Ericsson, and SAP, 19 September 2024; Fried, Ina, "Scoop: Meta won't offer future multimodal AI models in EU," *Axios*, 17 July 2024; Stolton, Samuel and Gurman, Mark, "Apple won't roll out AI tech in EU market over regulatory concerns," *Bloomberg*, 21 June 2024; Davies, Pascale, "OpenAI launches first AI agent Operator but it won't be coming to Europe yet," *Euronews*, 24 January 2025.

While the media has tracked individual instances of delays and withheld releases of LLMs – e.g. Anthropic's non-release of Claude 2 and 2.1 in the EU,[2] or Meta's exclusion of EU users from its multimodal Llama models[3] – we are not aware of anyone having assessed LLM release data systematically as we do here.

By doing so, we hope to inform policymakers considering the design and implementation of AI-relevant regulation. Overly stringent regulation could cause release delays or even regulatory flight.[4] At the same time, advances in AI capabilities may increase risks stemming from AI.[5] If policymakers weaken regulation to avoid release delays, the resulting framework may be inadequate for mitigating those risks.

Most imminently, we think this report could have implications for the implementation of the EU's AI Act – which will commence enforcement of its general-purpose AI provisions in August 2026.[6]

We note that, in the first half of 2026, constraints on European access to frontier AI were driven by trusted-access programs (GPT 5.5 Cyber, Claude Mythos, and GPT Rosalind) and US export controls (which led to the withdrawal of Claude Fable 5 and Mythos 5) – rather than regulation. Such releases are beyond the scope of our report, which focuses on public releases in the US.[7] Nonetheless these instances are noteworthy: It may be that US national security concerns, rather than European regulation, prove the more important barrier to European access to frontier AI in the coming years.

---

[2] Initially, Claude, powered by Claude 2 model, was launched in the US and UK only (Anthropic, "Claude 2," 11 July 2023). By the time Claude was available in the EU in May 2024, it was powered by Claude 3, and there was no option to switch to Claude 2 via the web app (Anthropic, "Claude is now available in Europe," 14 May 2024).

[3] Meta, "Llama 3.2 Acceptable Use Policy", *GitHub*, archived on 27 September 2024; Meta, "Llama 4 Acceptable Use Policy", GitHub, archived on 6 April 2025.

[4] By "regulatory flight" we mean the risk that companies relocate operations, withhold products, or reduce investment in jurisdictions with higher regulatory burdens. For discussion of how EU regulation may shape global AI markets, including the potential for such effects, see: Siegmann, Charlotte and Anderljung, Markus, "The Brussels Effect and AI: How EU Regulation Will Shape the Global AI Market," *GovAI*, 2022.

[5] The second International AI Safety Report, published in February 2026, synthesizes the current evidence on risks from general-purpose AI, including risks from malicious use (e.g. cyberattacks, disinformation, biological threats), technical failures (e.g. unreliable behavior, loss of control), and systemic risks (e.g. dependence on a small number of providers, cascading failures). See Bengio, Yoshua et al., "International AI Safety Report," February 2026, in particular Chapters 3–5.

[6] The general-purpose AI (GPAI) model provisions applied from 2 August 2025, though the AI Office's enforcement powers only apply from 2 August 2026. More broadly, the AI Act (Regulation (EU) 2024/1689) entered into force on 1 August 2024 but applies on a staggered timeline: provisions on prohibited AI practices and AI literacy obligations applied from 2 February 2025, while requirements for high-risk AI systems were originally due to apply from 2 August 2026 (standalone systems, Annex III) and 2 August 2027 (product-embedded systems, Annex I). On 7 May 2026, the Parliament and Council reached a provisional agreement on the Commission's Digital Omnibus on AI (proposed 19 November 2025), postponing these high-risk deadlines to 2 December 2027 and 2 August 2028 respectively, without affecting the GPAI obligations or the 2 August 2026 enforcement date; formal adoption by the Council is expected before 2 August 2026. See European Commission, "AI Act"; Council of the EU, "Artificial Intelligence: Council and Parliament agree to simplify and streamline rules," 7 May 2026.

[7] Project Glasswing – Anthropic's invite-only trusted-access program for Claude Mythos – is the closest borderline case (Anthropic, "Project Glasswing: Securing critical software for the AI era," 7 April 2026). European organizations were initially absent from the consortium of organizations given access under Glasswing, which some may read as delayed access. We treat it as out of scope because it was a restricted, invite-only release rather than a public one (see Methodology).

2. Delays and Non-Releases of LLMs in the EU and the UK

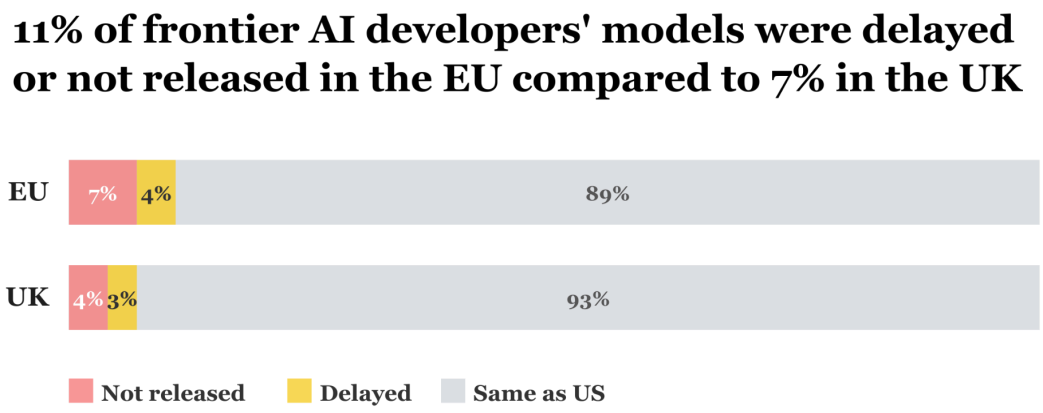

Figure 1: We analyzed a dataset of 375 LLM releases from Meta, Google, OpenAI, and Anthropic (06/2018 – 05/2026) to compare US release dates with EU and UK release dates. We show here the overall percentage of LLMs that were not released, delayed, or released simultaneously within the EU and UK, relative to the date of their US release. All numbers are rounded to the nearest whole number.

Overall, relative to the US, we find that ~7% of LLM releases were delayed or not released to the UK, and 11% were delayed or not released to the EU.

In the UK, ~3% of LLM releases were delayed, with a median delay of 15 days, and ~4% were never released. In the EU, ~4% of LLM releases were delayed, with a median delay of 53 days, and ~7% were never released. In~6% of cases, the UK received access to a model earlier than the EU, or was able to access a model that was not released in the EU. Nearly half of these cases were attributable to Meta.

What about the most performant models? Based on data from the Epoch Capabilities Index, the UK saw no delays in accessing the highest performing model at any one point in time. The EU saw a 71-day delay in accessing the web-app release of Anthropic's Claude 3 Opus, but no delay in gaining API access.[8]

[8] Epoch AI, "Capabilities Index," accessed 1 April 2026. The index combines scores from many different AI benchmarks into a single "general capability" scale, which can be used to proxy the most generally capable model since March 2023. During this period, we found there were 14 models in our dataset that were the highest rated at any one time, and of these only one (Claude 3 Opus) was delayed in its release to the UK or EU.

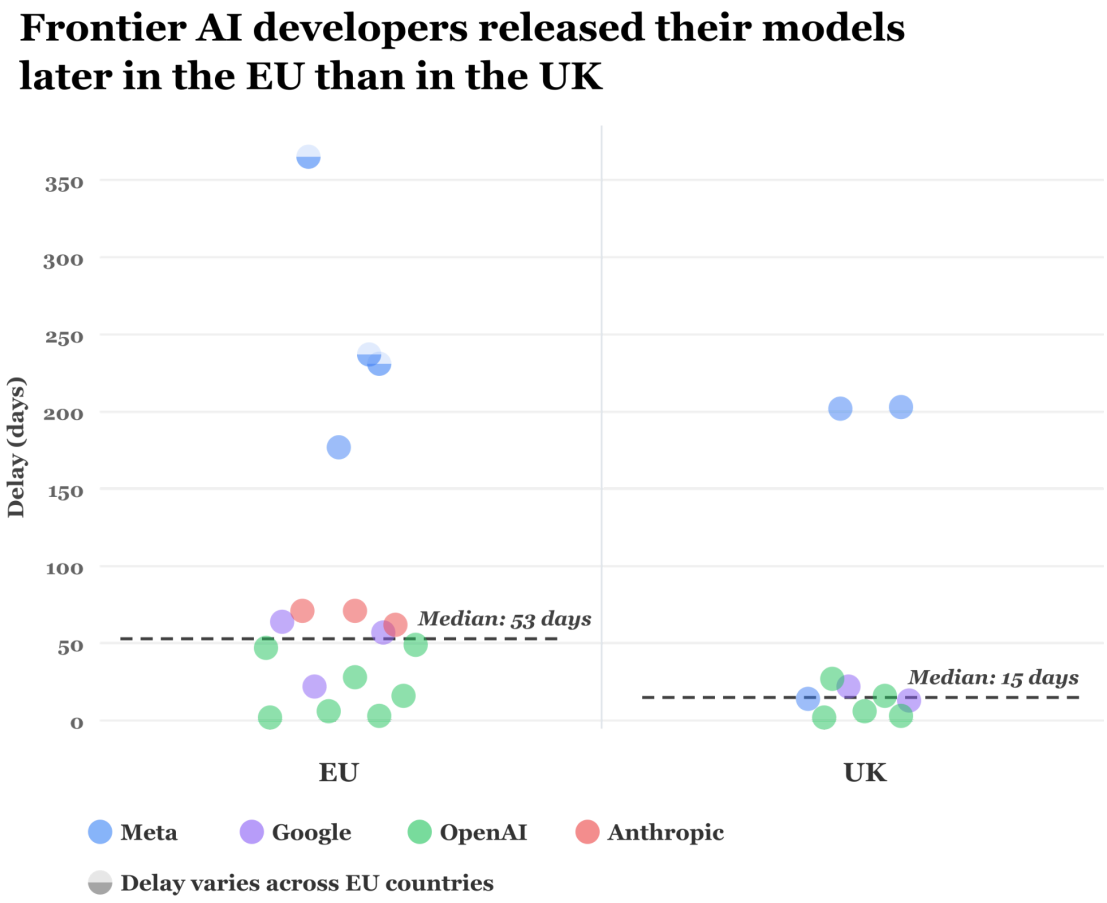

Figure 2: We analyzed a dataset of 375 LLM releases from Meta, Google, OpenAI, and Anthropic (06/2018 – 05/2026) to compare US release dates with EU and UK release dates. Here we show all the instances of a model being delayed to the EU and UK, and median delay per region. Each dot represents an LLM release, which is color-coded according to the company. We use a half-dot where there is variation between EU countries in the release date (two instances).

Delays and non-releases are not unique to LLMs. AI companies have also delayed the release of other AI and digital services:

- Google released AI Overviews in the UK three months later than in the US and roughly one year later in the EU.[9] The UK delay appears to be due to a staged international rollout,[10] while the greater EU delay appears to be due to regulatory uncertainty.[11]

[9] 10-month delay for Austria, Belgium, Germany, Ireland, Italy, Poland, Portugal, Spain (Budaraju, Hema, "We're bringing the helpfulness of AI Overviews to more countries in Europe," *Google*, 25 March 2025); 12-month delay for the rest of the EU, France excluded (Budaraju, Hema, "AI Overviews are now available in over 200 countries and territories, and more than 40 languages," *Google*, 20 May 2025); not released in France as of late June 2026, Google, "AI Overviews availability", (archived on 26 April 2026), see also Lainé, Linda, "AI Overviews (Google) n'est toujours pas déployé en France, mais les discussions avancent", *L'Echo Touristique*, 24 April 2026.
[10] Budaraju, Hema, "New ways to connect to the web with AI Overviews," *Google*, 15 August 2024.
[11] Kroet, Cynthia, "Google's AI Feature on Hold in Most EU Member States Due to 'Strict Rules'," *Euronews*, 1 April 2025.

- Meta withheld Threads from the EU for five months, apparently to ensure compliance with Digital Markets Act rules which restrict data-sharing between apps (Instagram and Threads in this case).[12]
- Apple declined to roll out Apple Intelligence, iPhone Mirroring and SharePlay Screen Sharing in the EU, citing regulatory uncertainties under the Digital Markets Act's interoperability requirements.[13]

Since we could not find a systematic review of release data for these services, we cannot tell if delays and non-releases are more severe for LLMs relative to other comparable products.

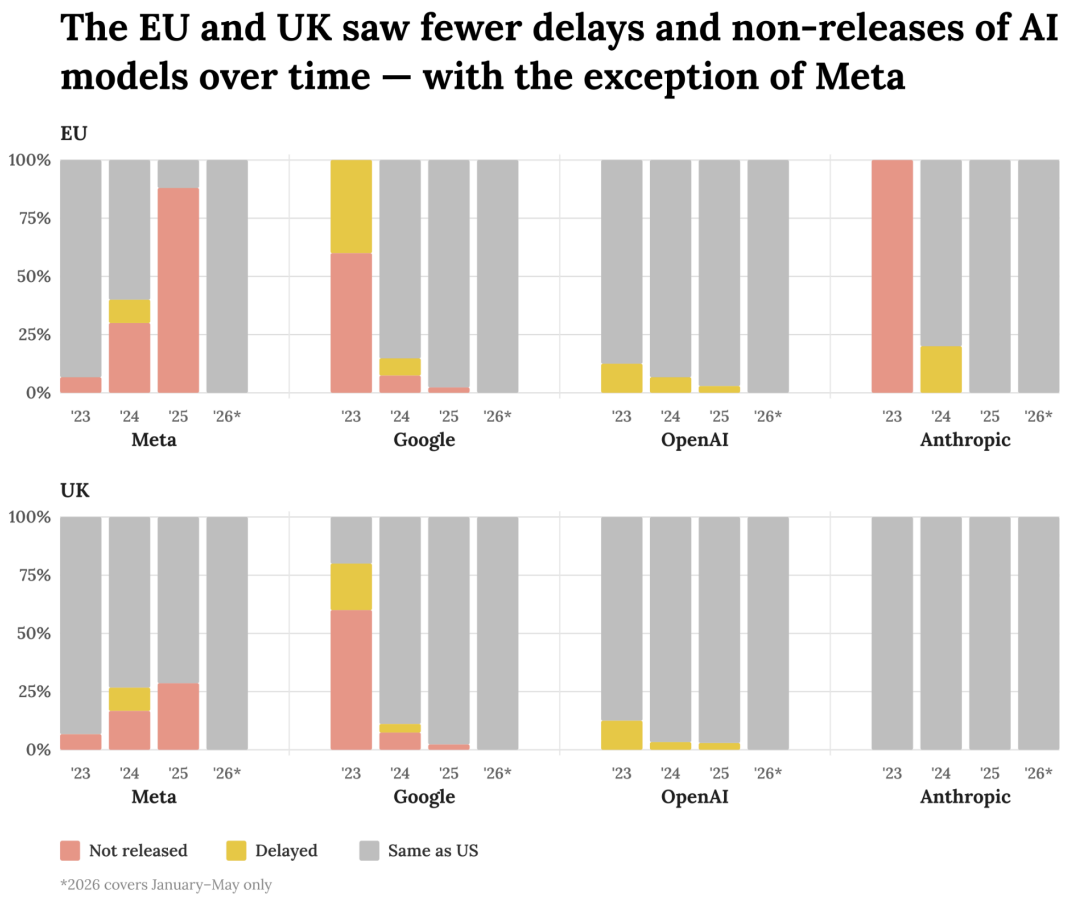

Figure 3: We analyzed a dataset of 375 LLM releases from Meta, Google, OpenAI, and Anthropic (06/2018 – 05/2026) to compare US release dates with EU and UK release dates. Here we show what percentage of LLM releases were not released, delayed, or released simultaneously relative to the US for each company per year during Jan 2023 – May 2026. We exclude releases between 2018 – 2022 because Anthropic and Google did not release models during this period.

[12] Dent, Steve, "Meta's Twitter Rival Threads Won't Initially Launch in the EU," *Engadget*, 5 July 2023; Mehta, Ivan, "Threads is finally available to users in the EU," *TechCrunch*, 14 December 2023.
[13] Stolton, Samuel and Gurman, Mark, "Apple Won't Roll Out AI Tech in EU Market over Regulatory Concerns," *Bloomberg*, 21 June 2024.

Of the four companies in the dataset, Meta has the highest share of LLMs that were delayed or withheld in the EU (~26%) and the UK (~15%). These mostly concern models used to power its AI assistant in Meta AI Glasses and its various digital platforms (WhatsApp, Facebook, Messenger, Instagram, and Meta.AI), or open source releases of multimodal Llama models.

Google delayed or withheld releases in ~11% of cases in the EU and 9% in the UK. Most of these (~60%) are for API releases of PaLM 2, Gemini 1.0, and 1.5 during 2023-2024.

OpenAI is the company with the lowest share of its LLMs delayed or not released in the EU (~4%). It is also the only company that has no instances of non-releases in either jurisdiction. Most OpenAI delays are relatively short (< 1 month), with a few longer delays for the EU release of GPT-4o with real-time audio and video input (49 days) and Operator (47 days).

Anthropic is the only company with no delays or non-releases in the UK. Though several models were either delayed or not released in the EU (Claude 2, 2.1, and 3) in 2023 and 2024, it is the only company in our dataset with no delays and non-releases in 2025 for the UK and EU.

Between 2023-2025, the share of delayed or non-released models per year has decreased over time both in the UK (from ~18% of total releases in 2023 to ~4% in 2025) and EU (from ~28% in 2023 to ~8% in 2025). Meta is the only company that has a reversed trend: The share of Meta's delayed or non-released models per year grew from ~7% to ~25% in the UK and from ~7% to ~88% in the EU over this period (see Figure 3).

For OpenAI, Google, and Anthropic, these trends may partly reflect companies transitioning from research to commercial organizations, with more mature compliance functions and greater incentive to deploy broadly. It could also reflect greater regulatory certainty as EU and UK regulators clarify how data protection laws apply to frontier LLM products.

3. Reasons for Delays and Non-Releases

Here we assess cases of delays and non-releases from the dataset and label the reason as either regulatory barriers or other factors, based on our interpretation of public sources.

We assign "regulatory factors" if regulation appears to be the primary reason behind the delay or non-release. In one case, a regulator explicitly prohibited the release of a model (the EU release of Google's Bard); in most cases, regulation appears to cause a delay due to the additional time needed to assure compliance or resolve legal uncertainty. This sometimes leads to a non-release where a model is deprecated in the meantime. Some of Google's free-tier API releases appear to have been withheld because EU data protection rules prohibited the use of user inputs for model training, making it commercially undesirable rather than legally prohibited to release the model.

We assign "other factors" where a non-regulatory factor appears to be the primary reason behind the delay or non-release. These include staged international rollouts that may reflect compute constraints (e.g. some of OpenAI's earlier releases) or caution about new high-risk capabilities (e.g. OpenAI's early agentic models). Language adaptation contributed to some delays, e.g. Meta's EU rollout of its AI assistant for Meta Glasses in certain EU countries.

Of the 68 examples of delays and non-releases in the dataset, we tentatively attribute regulatory factors as the primary, though not sole, cause of 56. Data protection regulation is the only type of regulation that we judge to have led to delays or non-releases for the model releases covered in our dataset (**Box 1**). We considered (among others) competition, copyright, and online safety regulations in the EU and UK, and the AI Act in the EU – but did not find sufficient evidence of them blocking releases of the models covered in our dataset.

Box 1: Data Protection Regulation as the Main Regulatory Barrier

EU and UK data protection regulation and LLMs

The General Data Protection Regulation (GDPR) in the EU, and the UK GDPR and Data Protection Act 2018 in the UK, govern how organizations process personal data, requiring a lawful basis, transparency, purpose limitation, and granting individuals rights to access, correct, and erase their data.[14]

Because LLMs are trained and operated on large volumes of personal data, they raise compliance issues such as establishing a lawful basis for training on personal data, meeting transparency and data-minimization requirements and satisfying individual rights, such as erasure, once personal data has been incorporated into model parameters.[15]

Examples of data protection regulation creating barriers in our dataset

We find 56 examples of delays and non-releases that appear primarily due to regulatory barriers. Within these examples, we find evidence of EU and UK data protection regulation appearing to pose barriers in three main ways:

- **Lawful basis for training on personal data.** Training a model on personal data requires a lawful basis under the GDPR, such as consent or legitimate interest, and this is legally contested when data that people shared for other reasons – for example social media posts – is reused to train AI.[16] EU and UK regulators

[14] General Data Protection Regulation (Regulation (EU) 2016/679); the UK GDPR (the retained EU law version of Regulation (EU) 2016/679, as amended by the Data Protection, Privacy and Electronic Communications (Amendments etc.) (EU Exit) Regulations 2019); and the Data Protection Act 2018.

[15] Dewitte, Pierre, "AI Meets the GDPR: Navigating the Impact of Data Protection on AI Systems," in Nathalie A. Smuha (ed.), The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence, Cambridge University Press, 2025, Chapter 7, pp. 133–157. See also Feretzakis, Georgios et al., "GDPR and Large Language Models: Technical and Legal Obstacles," Future Internet 17, no. 4 (2025): 151.

[16] GDPR, Article 6(1) (lawful bases for processing, in particular consent under Article 6(1)(a) and legitimate interests under Article 6(1)(f)). On the contested application of legitimate interest to AI model training, see EDPB, "Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models," 17 December 2024.

scrutinized Meta's use of first-party Facebook and Instagram data to train its models, which appears to have contributed to delays and non-releases.[17]

- **Stricter rules for sensitive (special-category) data.** The GDPR imposes stricter conditions on "special category" data – information revealing characteristics such as health, ethnicity, or religion (Article 9), which are more likely to be contained in images, voice, and real-time video data.[18] This may help explain why Meta excluded EU users from its image-capable (multimodal) Llama models and withheld image-input features from its assistant, and why OpenAI's real-time audio and video input for GPT-4o was delayed in the EU.
- **Insufficient compliance documentation.** Before deploying higher-risk processing, the GDPR requires a provider to assess and document its compliance, in particular through a data protection impact assessment (DPIA) under Article 35, and a regulator can delay a launch where this documentation is missing or inadequate.[19] The Irish DPC blocked the EU launch of Google's Bard after Google had not provided the required impact assessment; Google proceeded only after addressing the regulator's concerns, including new transparency information, user controls, and a data-collection opt-out.[20]

The EU appears to pose greater data protection regulatory barriers than the UK

Although the EU and UK share broadly the same data protection law, regulatory barriers appear to have been higher in the EU.[21] We see two main reasons for this:

- **EU enforcement appears to be more assertive, which may make companies cautious even without engaging with a regulator.** Italy's Garante temporarily banned ChatGPT in 2023, and France's CNIL penalized Clearview AI for scraping images from the web to train facial recognition software, including a €5.2 million fine for non-compliance.[22] The Irish DPC separately opened an investigation into Meta's use of first-party Facebook and Instagram data for AI training.[23]

[17] ICO, "Statement in Response to Meta's Plans to Train Generative AI with User Data," 14 June 2024; DPC, "The DPC's Engagement with Meta on AI," 14 June 2024.
[18] GDPR, Article 9 (processing of special categories of personal data). On the principle that processing "liable to indirectly reveal" protected characteristics falls within Article 9, see Case C-184/20, OT v Vyriausioji tarnybinės etikos komisija, EU:C:2022:601.
[19] GDPR, Article 35
[20] Graham Doyle, DPC Deputy Commissioner, quoted in "Google Delays EU Launch of Its AI Chatbot After Privacy Regulator Raises Concerns," TechCrunch, 13 June 2023; see also Lomas, Natasha, "Google's Bard and Other AI Chatbots Remain Under Privacy Watch in the EU," TechCrunch, 13 July 2023.
[21] Dewitte, "AI Meets the GDPR"; see also Arnal, Judith, "AI at Risk in the EU: It's Not Regulation, It's Implementation," European Journal of Risk Regulation 16 (2025): 1168–1177 arguing that the EU's implementation of the GDPR and its fragmented governance architecture, rather than the rules themselves, hold back AI deployment, with Meta as the "paradigmatic" case.
[22] Italian Data Protection Authority (Garante), "Artificial Intelligence: Stop to ChatGPT by the Italian SA," 31 March 2023. On the Clearview penalty, CNIL (French Data Protection Authority), "Facial Recognition: the CNIL Imposes a Penalty Payment on CLEARVIEW AI," 10 May 2023 (€5.2 million for failure to comply with the order accompanying the CNIL's €20 million fine of 17 October 2022).
[23] DPC, "The DPC's Engagement with Meta on AI," 14 June 2024.

- **The EU tends to take longer to reach or clarify a regulatory position.** Responsibility for data protection is divided among national authorities and coordinated at EU level through the European Data Protection Board (EDPB), whereas the UK has a single regulator, the ICO. When Meta sought to train on first-party public data, the ICO reached a position in about three months, while in the EU the EDPB took months to issue an opinion that each national authority then had to interpret and apply individually, leaving the legal position less certain.[24]

Nine instances of delays or non-releases appear to be primarily driven by non-regulatory factors, such as compute constraints, language barriers, or product readiness concerns.

There are three instances of delays or non-releases we have labeled as "unclear", where we have not been able to find sufficient evidence in favor of regulatory or non-regulatory factors as the primary cause.

We have high confidence about the reasons for only five instances of a delay or non-release, due to a lack of strong public evidence about their causes. In most cases, there is no direct statement from the company about the reason for the model being delayed or not released into the EU or UK. For the cases where a reason is given, it tends to be at a very high level. And even where regulatory barriers are cited as the reason, this should not be taken at face value: In some cases it may be strategically valuable to cite regulation as the cause, to increase public pressure on regulators towards more permissive enforcement.

[24] On the UK timeline, ICO, "Statement in Response to Meta's Plans to Train Generative AI with User Data," 14 June 2024, and ICO, "Statement in Response to Meta's Announcement on User Data to Train AI," 13 September 2024 (a position reached in roughly three months); on the EU timeline, EDPB, "Opinion 28/2024...," 17 December 2024, issued following the Irish DPC's request of September 2024.

Table 1: LLM releases by Meta, Google, OpenAI, and Anthropic (06/2018 – 05/2026) that were delayed or withheld to the UK or EU relative to the US, and the primary cause for the delay or non-release.

*WA = WhatsApp, ME = Messenger, FB = Facebook, IG = Instagram.

Model	Mode	US release	UK delay	EU delay	Main factor	Confidence
META						
BlenderBot 3.0 175B Public Preview	Consumer App	05/08/2022	Not released	Not released	Other (both)[25]	Medium
Llama 2 (unspecified model)	Consumer App (WA, ME, IG)	06/12/2023[26]	Not released	Not released	Other (both)[27]	Medium
Llama 3 (unspecified model)	Consumer App (WA, ME, FB, Meta.AI)	18/04/2024	Not released	Not released	Reg. (both)[28]	Medium
Llama (unspecif. model, no image input)	Consumer App (Meta Glasses)	23/04/2024	202 days	237 & 365 days[29]	Reg. (both)[30]	Medium

[25] Meta AI, "BlenderBot 3: A 175B Parameter, Publicly Available Chatbot That Improves Its Skills and Safety Over Time," 5 August 2022. BlenderBot 3 launched as a US-only public research demo designed to collect conversational feedback data. A Meta researcher noted they were still working on expanding outside the US (Weston, Jason [@jaseweston], "Hi! What's unique in this project...," *X*, 5 August 2022, archived on 26 April 2026), but this never happened, likely because the project was superseded by newer models. Confidence is medium because while there is no evidence of regulatory barriers, we cannot rule out that compliance costs contributed to the decision to expand to the UK and EU later.

[26] Beta version of Meta AI assistant powered by Llama 2 was released in the US on September 27, 2023 (Meta AI, "Introducing New AI Experiences Across Our Family of Apps and Devices," 27 September 2023). However, we do not include this as a release date in the table, because it seems that generally, Meta's beta releases seem to be private, e.g. limited number of sign-ups, and we recorded the data for public releases only.

[27] Meta, "What's New Across Our AI Experiences," 6 December 2023. Meta made its AI assistant powered by Llama 2 available to anyone in the US on WhatsApp, Messenger, and Instagram. The first international expansion did not occur until April 2024, but already with a newer Llama (Meta, "Meet Your New Assistant: Meta AI, Built With Llama 3," 18 April 2024), suggesting the UK and EU exclusion at this stage reflects a staged product rollout rather than jurisdiction-specific regulatory barriers. Confidence is medium because while there is no evidence of regulatory barriers, we cannot rule out that compliance costs contributed to the decision to expand to the UK and EU later.

[28] Meta, "Meet Your New Assistant: Meta AI, Built With Llama 3," 18 April 2024. Meta rolled out its AI assistant powered by Llama 3 to a limited set of countries including the US, Canada, Australia, and New Zealand, but excluded the UK and EU. Meta had informed the Irish Data Protection Commission (DPC) in March 2024 of its plans to train models on EU user data and had not yet received a response (DPC, "DPC Statement on Meta AI," 21 May 2025). At this point neither the UK's Information Commissioner's Office (ICO) nor the Irish DPC had clarified the legality of using first-party platform data to train AI models, creating regulatory uncertainty that plausibly contributed to the exclusion of both jurisdictions. Meta could in principle have released a version not trained on EU data, but later described this option as offering only a "second-rate experience" (Fratta, Stefano, "Update to Building AI Technology for Europeans in a Transparent and Responsible Way," *Meta*, 14 June 2024). Confidence is medium because many other countries were also excluded at this stage, and the non-release may partly reflect a commercial decision not to ship what Meta considered an inferior product rather than a hard regulatory barrier.

[29] Austria, Belgium, Denmark, Finland, Germany, and Sweden delayed 365 days; France, Ireland, Italy, and Spain delayed 237 days; all other EU countries did not sell Meta Glasses during the period covered in our dataset.

[30] Meta AI features launched on glasses in the US and Canada only (Meta, "New Ray-Ban Meta Smart Glasses Styles and Meta AI Updates," 23 April 2024); when the glasses eventually reached parts of Europe, they initially offered voice-only features (Meta, "AI Glasses Release Notes," for the UK - update for 11 November 2024, for the EU - update for 16 December 2024). Meta explicitly cited Europe's "complex regulatory system" as the reason for delay (Meta, "Meta AI Begins Roll Out on Ray-Ban Meta Glasses in France, Italy, Ireland and Spain," archived on 18 November 2024). Smart glasses pose distinct challenges under the General Data Protection Regulation (GDPR): The built-in camera can capture images of bystanders who have not consented to processing, potentially including data from which protected characteristics could be inferred under Article 9 of GDPR. Language capabilities might have been an additional factor in the EU rollout sequencing: The first EU countries to gain access were France, Italy, Spain, and Ireland, matching supported languages at the time. No other company has launched AI glasses with visual AI features in Europe, so there is no counterfactual to compare against. Confidence is medium because while Meta's own statement is strong evidence, we lack visibility into the specific compliance barriers Meta faced.

Model	Mode	US release	UK delay	EU delay	Main factor	Confidence
Llama 3.1 (unspecified model)	Consumer App (ME, IG, FB)	23/07/2024	Not released	Not released	Reg. (both)[31]	Medium
Llama 3.1 (70B)	Consumer App (WA, Meta.AI)	23/07/2024	Not released	Not released	Reg. (both)[31]	Medium
Llama 3.1 (405B)	Consumer App (WA, Meta.AI)	23/07/2024	Not released	Not released	Reg. (both)[31]	Medium
Llama 3.2 Vision (11B)	Open Source	25/09/2024	Same as US	Not released	Reg. (EU)[32]	Low
Llama 3.2 Vision (90B)	Open Source	25/09/2024	Same as US	Not released	Reg. (EU)[32]	Low
Llama 3.2 Vision Instruct (11B)	Open Source	25/09/2024	Same as US	Not released	Reg. (EU)[32]	Low
Llama 3.2 Vision Instruct (90B)	Open Source	25/09/2024	Same as US	Not released	Reg. (EU)[32]	Low
Llama 3.2 (unspecified model, image input)	Consumer App (WA, ME, FB, IG, Meta.AI)	25/09/2024	Not released	Not released	Reg. (both)[33]	Low

[31] Meta expanded its AI assistant from 15 to 22 countries, excluding the UK and EU (Meta, "Meta AI Is Now Multilingual, More Creative and Smarter," archived on 23 July 2024). Both the UK Information Commissioner's Office (ICO, "Statement in Response to Meta's Plans to Train Generative AI with User Data," 14 June 2024) and the Irish Data Protection Commission (DPC, "The DPC's Engagement with Meta on AI," 14 June 2024) had requested Meta pause its plans to train models on UK and EU user data just one month earlier, making regulatory factors the most likely primary cause of the continued exclusion. Meta could in principle have released a version not trained on local data, but described this as a "second-rate experience" it was unwilling to offer (see footnote 28). Confidence is medium because many other countries also lacked access at this point, and the non-release may partly reflect a commercial decision not to ship what Meta considered an inferior product rather than a hard regulatory barrier.

[32] The Llama 3.2 multimodal models were withheld from EU individuals and companies under Meta's own Community License terms (Meta, "Llama 3.2 Acceptable Use Policy", *GitHub*, archived on 27 September 2024), while text-only models remained available globally. Meta had already told Axios in July 2024 that it would not offer future multimodal models in the EU because of regulatory uncertainty (Fried, Ina, "Scoop: Meta Won't Offer Future Multimodal AI Models in EU," *Axios*, 17 July 2024). The UK was not excluded, possibly because the UK Information Commissioner's Office (ICO) had indicated it would not block Meta's training plans (ICO, "Statement in Response to Meta's Announcement on User Data to Train AI," 13 September 2024), while the equivalent EU process remained unresolved. Neither the UK nor the EU received the multimodal models for the AI assistant on Meta's platforms (see footnote 33), suggesting different considerations apply to the open source release and the assistant release. Confidence is low because the UK/EU divergence could reflect differences in regulatory posture between the UK and EU, or a strategic decision by Meta to pressure EU regulators.

[33] This is the first release in which Meta rolled out a multimodal Llama model in its AI assistant on its digital platforms; neither the UK nor the EU received it, and neither has as of June 2026. This contrasts with the open source multimodal release (footnote 32), where the UK was included but the EU was not. This is likely due to a similar barrier to that discussed in footnote 28, relating to Meta's use of first-party platform data for training. The barrier may be harder for image-capable features, which raise more complex privacy and transparency questions than text chat, particularly where photographs of identifiable people are involved. The non-release in the UK as well as the EU, despite the UK Information Commissioner's Office having indicated it would not block Meta's training plans, suggests that the Meta AI assistant release involves additional compliance considerations beyond the legality of the training data itself. Confidence is low because Meta has not publicly cited regulation as the reason for this non-release, and we cannot distinguish between a regulatory barrier and a commercial decision not to ship an inferior product.

Model	Mode	US release	UK delay	EU delay	Main factor	Confidence
Llama 3.2 (unspecified model, no image input)	Consumer App (ME, IG, FB, Meta.AI)	25/09/2024	14 days	177 days	Reg. (both)[34]	High (UK) Medium (EU)
Llama (unspecified model, image input)	Consumer App (Meta Glasses)	25/09/2024	203 days	231 days[35]	Reg. (both)[36]	Low
Llama 4 Scout	Open Source	05/04/2025	Same as US	Not released	Reg. (EU)[37]	Low
Llama 4 Scout Instruct	Open Source	05/04/2025	Same as US	Not released	Reg. (EU)[37]	Low
Llama 4 Maverick	Open Source	05/04/2025	Same as US	Not released	Reg. (EU)[37]	Low
Llama 4 Maverick Instruct	Open Source	05/04/2025	Same as US	Not released	Reg. (EU)[37]	Low
Llama 4 (unspecified model, image input)	Consumer App (WA, ME, IG, Meta.AI)	16/04/2025	Not released	Not released	Reg. (both)[38]	Low
Llama 4 (unspecified model, image input)	Consumer App (Meta.AI)	29/04/2025	Not released	Not released	Reg. (both)[38]	Low
GOOGLE						

[34] The Meta AI assistant powered by text-only Llama 3.2 reached the UK after 14 days (Meta, "Update to Meta AI Is Now Multilingual, More Creative and Smarter," 9 October 2024) and the EU after 177 days (Meta, "Europe, Meet Your Newest Assistant: Meta AI," 19 March 2025). The UK launch in October 2024 followed shortly after the UK Information Commissioner's Office (ICO) said Meta had changed its approach and that it would monitor resumed processing – though the ICO explicitly stated it had "not provided regulatory approval" (ICO, "Statement in Response to Meta's Announcement on User Data to Train AI," 13 September 2024). The EU launch did not come until March 2025, using a model not trained on first-party data from EU users (Anna Dack, Meta's innovation communications manager, quoted in *TechCrunch*, "Meta AI Is Finally Coming to the EU, but with Limitations," 20 March 2025) – just weeks before Llama 4 launched in the US in April. Meta could presumably have built and released such a model sooner than it did, though we do not know how long it took to prepare the release. Confidence is high for the UK, where the delay maps cleanly onto the ICO timeline, but medium for the EU. While Meta stated in the March release that the delay was caused by the EU's "complex regulatory system," we cannot rule out with confidence that it was due to strategic withholding or needing more time to potentially develop a model adapted to the EU market, or a combination of these.

[35] Austria, Belgium, Denmark, Finland, France, Germany, Ireland, Italy, Spain, and Sweden rollout was delayed by 231 days; all other EU countries did not sell Meta Glasses during this period.

[36] Meta AI visual features ("ask about what you see") launched on glasses in the UK in the week of 16 April, 2025, and in EU countries where the glasses are sold in the week of 14 May, 2025 (Meta, "AI Glasses Release Notes"). The delay relative to the US launch (September 2024) was substantial. Meta did not specifically address the visual features rollout sequencing, but had previously cited Europe's "complex regulatory system" as the reason for the delay to the underlying Meta AI rollout on glasses (see footnote 30). The fact that the AI model with image input capabilities was eventually released on glasses for the UK and EU but not on the Meta AI assistant (see footnote 33) – despite both using supposedly the same underlying model – could be due to compliance differences between the two release contexts. Confidence is low because we lack visibility into what specific compliance steps enabled such deployment for the AI glasses but not the Meta AI assistant.

[37] We speculate this is due to the same reasons set out in footnote 32.

[38] We speculate this is due to the same reasons set out in footnote 33.

Model	Mode	US release	UK delay	EU delay	Main factor	Confidence
PaLM 2 (unspecified model)	Consumer App	10/05/2023	Same as US	64 days	Reg. (EU)[39]	High
PaLM 2 - Text Bison 001	API	28/08/2023	Not released	Not released	Reg. (both)[40]	Medium
PaLM 2 - Chat Bison 001	API	28/08/2023	Not released	Not released	Reg. (both)[40]	Medium
Gemini 1.0 Pro (unspecified model)	Consumer App	06/12/2023	13 days	57 days	Reg. (both)[41]	Low
Gemini 1.0 Pro Vision	API	13/12/2023	Not released	Not released	Reg. (both)[40]	Medium
Gemini Ultra 1.0	Consumer App	08/02/2024	Same as US	N.R. in France	Unclear (EU)[42]	–
Gemini 1.5 Pro Public Preview	API (Free)	09/04/2024	Not released	Not released	Reg. (both)[40]	Medium
Gemini 1.5 Flash Public Preview	API	14/05/2024	Not released	Not released	Reg. (both)[40]	Medium

[39] Google simultaneously released its web app Bard powered by PaLM2 in the US, UK, and 181 other countries and territories in May 2023, excluding the EU, other European countries, and Canada (Google, "Where You Can Use Bard," archived on 11 May 2023). Google had intended to launch in the EU in mid-June 2023, but the Irish Data Protection Commission (DPC) blocked the rollout after Google failed to provide a data protection impact assessment or adequate documentation (Graham Doyle, DPC Deputy Commissioner, quoted in *TechCrunch*, "Google Delays EU Launch of Its AI Chatbot After Privacy Regulator Raises Concerns," 13 June 2023). Bard powered by PaLM2 launched in the EU approximately one month later, on July 13, 2023, after Google made changes including increased transparency, user controls, and an opt-out for data collection (Lomas, Natasha, "Google's Bard and Other AI Chatbots Remain Under Privacy Watch in the EU," *TechCrunch*, 13 July 2023). Google also agreed to provide the Irish DPC with a review report after three months of operation. The UK was not delayed, suggesting the UK's regulatory posture did not create friction at this stage. Confidence is high because the Irish DPC directly intervened to block the planned launch, and the delay corresponds to the time needed to satisfy the regulator's requirements.

[40] Up until May 2024, Google had a free-tier-only access to the PaLM and Gemini API that permitted Google to use submitted prompts and responses to "improve, and develop Google products and services and machine learning technologies" (Google, "Gemini API Additional Terms of Service," archived on 18 May 2024). This may conflict with the principle of purpose limitation under Article 5(1)(b) of the GDPR, which requires that personal data not be further processed in a manner incompatible with the purposes for which it was collected (see e.g. Dewitte, Pierre, "AI Meets the GDPR," *The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence*, 2025), and could be the reason why the European Economic Area (EEA) countries (EU included), UK, and Switzerland had a paid-tier-only access for some time, under which Google "doesn't use your prompts (including associated system instructions, cached content, and files such as images, videos, or documents) or responses to improve our products" (Google, "Gemini API Additional Terms of Service," archived on 18 May 2024). This meant, however, that the models that were exclusively free-tier were not accessible for EU and UK users. By the time Google ruled that paid-tier data handling terms applied to users in the EEA, UK, and Switzerland even on the free plan (Logan Kilpatrick, Google AI Studio Product Lead, "We just expanded the Gemini API free tier access...," *X*, 23 August 2024), the models listed in these rows had been deprecated or superseded, and so were never released in the UK or EU. The first free-tier Gemini API release to the UK and EU was Gemini 1.5 Pro Experimental, launched with a 22-day delay. Confidence is medium because Google's terms of service changes are consistent with our interpretation that GDPR purpose limitation drove the free-tier exclusion, but Google never publicly stated this was the reason, and we cannot rule out that commercial or operational factors also contributed to the decision not to release these specific models in the UK and EU.

[41] Google integrated Gemini 1.0 Pro into Bard on 6 December 2023, making it available in over 170 countries but excluding both the UK and EU (Google, "Where Bard with Gemini Pro is available", archived on 17 December 2023). This was a change from the earlier Bard launch, where the UK had launched simultaneously with the US. The UK received access after approximately 13 days (Williams, Holly, "Google rolls out upgrade to AI chatbot Bard across UK", *Independent*, 21 December 2023); the EU after 57 days, when Google expanded to over 40 countries in early February 2024 (Duboust, Oceane, "Google's latest AI model, Gemini Pro, is now available in Europe to compete with OpenAI's ChatGPT", *Euronews*, 2 February 2024). The delay was allegedly attributed by some to "regulatory hurdles" (Ahlgren, Linnea, "Google's Gemini AI Won't Be Available in Europe – For Now", *The Next Web*, 7 December 2023), but no source specifies which regulator or regulation caused the delay for either jurisdiction. Confidence is low because while regulatory factors are plausible, we lack direct evidence of a specific regulatory intervention.

[42] Gemini Advanced powered by Gemini Ultra 1.0 launched in 150+ countries but excluded France; the rest of the EU and the UK received access simultaneously with the US (Google, "Upgrade to Gemini Advanced", archived on 8 February 2024). Google did not publicly explain the exclusion, and we have been unable to identify a sufficiently credible reason.

Model	Mode	US release	UK delay	EU delay	Main factor	Confidence
Gemini 1.5 Pro Experimental	API	01/08/2024	22 days	22 days	Reg. (both)[40]	Medium
Gemini 2.0 Flash Preview Image Generation	API	07/05/2025	Not released	Not released	Unclear (both)[43]	–
OPENAI						
GPT-3.5 Turbo Alpha	Consumer App	10/02/2023	1 day	1 day	Other (both)[44]	High
GPT-4 (image input)	Consumer App	25/09/2023	16 days	16 days	Reg. (both)[45]	Low
GPT-4o (audio input and output)	Consumer App	24/09/2024	6 days	28 days	Reg. (both)[46]	Medium
GPT-4o (real-time audio and video input)	Consumer App	12/12/2024	Same as US	49 days	Reg. (EU)[47]	Low

[43] Google launched image generation capabilities in preview for Gemini 2.0 Flash on 7 May 2025 but did not release the feature in the UK or EU (Google, "Gemini models," archived on 8 May 2025, stating that the model is "not currently supported in a number of countries in Europe, Middle East & Africa"). Google did not publicly explain the exclusion, and we have been unable to identify a sufficiently credible reason.

[44] The GPT-3.5 Turbo alpha model became available to US-based ChatGPT paid subscribers on 10 February 2023 according to Central European Time (CET) (OpenAI, "ChatGPT Release Notes"), and non-US-based users started gaining access to the paid subscription, and thus, the GPT-3.5 Turbo model from 11 February 2023 according to CET. The delay therefore likely reflects the operational requirements of launching a paid product across multiple jurisdictions, such as configuring international payment processing and tax compliance.

[45] OpenAI announced image input capabilities for GPT-4 on 25 September 2023, rolling them out to paid subscribers (OpenAI, "ChatGPT Can Now See, Hear, and Speak", 25 September 2023). The UK and EU were specifically excluded at launch – OpenAI's help page at the time stated the feature was "Not yet available in the UK and EU" – while access was seemingly granted simultaneously to users in many other jurisdictions (OpenAI, "Image inputs for ChatGPT - FAQ," archived on 28 September 2023). UK and EU paid subscribers received the feature on 11 October 2023, 16 days later (OpenAI Developer Community, "Imagine input for UK and EU ChatGPT Plus users has arrived", 11 October 2023). OpenAI did not publicly explain the exclusion. We speculate that the delay reflects additional GDPR-related compliance review for the image input modality, given the simultaneous UK + EU release pattern matches other GDPR-jurisdiction-specific delays in the dataset, and because image data raises distinct compliance questions under GDPR Article 9. Confidence is low because OpenAI has not cited regulation as the reason and we cannot rule out non-regulatory explanations such as additional internal safety review for the new modality before extending it to jurisdictions where OpenAI faced active regulatory engagement.

[46] OpenAI launched audio input and output capabilities for GPT-4o for ChatGPT paid subscribers on 24 September 2024, excluding the EU, UK, Switzerland, Iceland, Norway, and Liechtenstein. The UK received access after 6 days (OpenAI, "Voice mode FAQ," archived on 2 October 2024); the EU after 28 days (OpenAI [@OpenAI], "Well yes. All Plus users in the EU...", *X*, 22 October 2024). OpenAI commented that "some parts of the world require additional external reviews before launching new products" to "ensure the feature aligns with local requirements" (Barlow, Graham, "OpenAI's Advanced Voice Mode Is Unavailable in the EU, and Now We Might Know Why," *TechRadar*, 3 October 2024). We speculate the reviews related to GDPR compliance requirements for processing voice data, which may constitute special category biometric data under Article 9 when used to uniquely identify individuals (Lawne, Richard, "Biometrics in the EU: Navigating the GDPR, AI Act," *IAPP*, 23 April 2025). Confidence is medium because OpenAI's language strongly implies regulatory factors but the specific regulatory barrier is uncertain.

[47] On 12 December 2024, OpenAI released a model with real-time audio and video input in the ChatGPT mobile apps for paid subscribers, but the EU, Switzerland, Iceland, Norway, and Liechtenstein were specifically excluded (OpenAI, "ChatGPT - Release Notes", 12 December 2024). We speculate that the delay reflects additional GDPR-related compliance review for the new visual modality added on top of voice – both of which raise distinct compliance questions, voice as potential special-category biometric data under Article 9 and image/video as data that may indirectly reveal protected characteristics. Confidence is low because OpenAI has not cited regulation as the reason and the specific compliance barrier is not identified.

Model	Mode	US release	UK delay	EU delay	Main factor	Confidence
GPT-4o Computer Use Agent ("Operator")	Consumer App	25/01/2025	27 days	47 days	Other (UK)[48]	High
					Reg. (EU)[49]	Low
o3 Deep Research (unspecified model)	Consumer App	02/02/2025	3 days	3 days	Other (both)[50]	Medium
ChatGPT agent	Consumer App	17/07/2025	Same as US	6 days	Reg. (EU)[51]	Low
ANTHROPIC						

[48] OpenAI initially released Operator powered by GPT-4o Computer Use Agent model in the US only, framing it as a "research preview" with an intentionally cautious rollout (OpenAI, "Introducing Operator," archived on 23 January 2025). OpenAI noted that as one of its first agentic models with the ability to take actions in a browser, Operator "brings new risks and challenges to address," including prompt injection attacks from third-party websites, irreversible model mistakes, and potential misuse (OpenAI, "Computer-Using Agent," archived on 23 January 2025). The UK received access after 27 days, on February 21, 2025, as part of the first broad international expansion alongside other non-EU markets (OpenAI [@OpenAI], "Operator is now rolling out to Pro users in Australia...," X, 21 February 2025). Compute constraints are a plausible alternative or additional explanation for the staggered rollout: Operator launched only to ChatGPT's $200-per-month Pro tier in the US, and as a browser-using agent it is compute-intensive, running a hosted browser alongside repeated model calls for each task. OpenAI was also openly capacity-constrained in this period – a week after the UK launch, CEO Sam Altman said the company was "out of GPUs" and was staggering the rollout of GPT-4.5 by subscription tier, adding "tens of thousands of GPUs" before widening access (Wiggers, Kyle, "OpenAI CEO Sam Altman Says the Company Is 'Out of GPUs'," TechCrunch, 27 February 2025). The staggered rollout therefore likely reflects a combination of these novel safety concerns and compute limits rather than UK-specific regulatory barriers.

[49] The EU, together with Switzerland, Iceland, Norway, and Liechtenstein, received access to Operator powered by GPT-4o Computer Use Agent model after 47 days, on 13 March 2025 (OpenAI [@OpenAI], "Operator Update," *X*, 13 March 2025), 20 days later than the UK and other non-EU markets. At launch, OpenAI CEO Sam Altman stated that Operator would reach other countries soon but that "Europe will, unfortunately, take a while," without elaboration (Davies, Pascale, "OpenAI Launches First AI Agent Operator but It Won't Be Coming to Europe Yet," *Euronews*, 24 January 2025). OpenAI did not publicly specify the reason for the additional EU delay. We speculate this could be due to the EU posing greater compliance barriers than the UK for novel AI products, given its more complex institutional structure and reputation for more aggressive enforcement of data protection regulation. Confidence is low because Altman's language implies regulatory factors (assuming that he used "Europe" to talk about the EU specifically, not the entire continent), but it is not clear what the specific barrier might be.

[50] Deep Research powered by a version of o3 launched for paid users in the US, with OpenAI stating it was "still working on bringing access to users in the United Kingdom, Switzerland, and the European Economic Area" (OpenAI, "Introducing Deep Research," 2 February 2025, archived on 3 February 2025). Access was extended to those regions three days later. In the same blog post, OpenAI noted that the feature is "very compute intensive". OpenAI may have delayed release to the EU and UK to help manage compute demand, staggering rollout per region. A three-day delay that was identical in the UK and EU fits a capacity-limited rollout better than a regulatory barrier. Confidence is medium because OpenAI did not explicitly link its compute-intensity statement to the decision to delay release to the EU and UK.

[51] ChatGPT agent model launched globally for paid users but explicitly excluded the European Economic Area (EEA; EU included) and Switzerland; the UK was not excluded (OpenAI, "Introducing ChatGPT Agent: Bridging Research and Action," 17 July 2025). OpenAI announced on 23 July that paid users in the EEA and Switzerland had been granted access (OpenAI [@OpenAI], "Pro users in the European Economic Area and Switzerland...," *X*, 23 July 2025). OpenAI did not explain the exclusion. We think regulatory factors could be involved because it is the EEA and Switzerland specifically, rather than the EEA and UK together, that were excluded, a pattern identical to the earlier Operator launch. The short delay could be additional compliance checks required for EU/EEA release. However, confidence is low because OpenAI has not cited regulation as the reason and we could not identify a specific regulatory barrier that seems likely to be the cause.

Model	Mode	US release	UK delay	EU delay	Main factor	Confidence
Claude 2	Consumer App	11/07/2023	Same as US	Not released	Reg. (EU)[52]	Low
Claude 2.1 70k context window	Consumer App	21/11/2023	Same as US	Not released	Reg. (EU)[52]	Low
Claude 2.1 200k context window	Consumer App	21/11/2023	Same as US	Not released	Reg. (EU)[52]	Low
Claude 3 Sonnet	Consumer App	04/03/2024	Same as US	71 days	Reg. (EU)[52]	Low
Claude 3 Opus	Consumer App	04/03/2024	Same as US	71 days	Reg. (EU)[52]	Low
Claude 3 Haiku	Consumer App	13/03/2024	Same as US	62 days	Reg. (EU)[52]	Low

Meta

Meta has the most instances of delays and non-releases of the four companies in the dataset. Most of these appear to be driven by regulatory factors. In particular, a number of delays/non-releases appear to be driven by regulatory uncertainty concerning Meta's use of first-party Facebook and Instagram data (data collected from users) for model training.

Two early cases – BlenderBot 3 (2022) and the Llama-2-powered AI assistant (2023) – appear to be driven by non-regulatory factors. BlenderBot 3 launched as a US-only public research demo and was never expanded internationally, likely because it was superseded by newer models.[53] The Llama-2-powered AI assistant was released to the US only as a beta test.[54]

[52] Claude web app powered by Claude 2 model launched in the US and UK only (Anthropic, "Claude 2", 11 July 2023). Anthropic later expanded the Claude-2-powered web app to 93 other countries but still excluded the EU, some other European and non-European countries (Anthropic, "Supported countries and regions: Claude.ai", archived on 14 October 2023), saying it was "working hard to responsibly expand availability" but giving no specific reason for the EU exclusion (Anthropic [@AnthropicAI], "We're working hard to responsibly expand availability...", *X*, 16 October 2023). At the time, Anthropic had no EU legal entity or GDPR-specific data processing framework. Anthropic was also a relatively young organization, founded in 2021. Claude powered by Claude 2 was Anthropic's first consumer-facing product, and the company may simply have lacked the resources to prioritise EU compliance alongside a global launch, with Claude powered by Claude 2.1 (Rows 39-40) following the same pattern. The eventual EU launch on 14 May 2024 (Rows 41-43; Anthropic, "Claude is now available in Europe", 14 May 2024) was preceded by structural preparation. On 22 March 2024, Anthropic incorporated Anthropic Ireland, Limited to act as its EU data controller under GDPR (SoloCheck, "Anthropic Ireland, Limited," archived on 8 May 2026). Around the same time of the eventual EU launch of the Claude web app, Anthropic also updated its Usage Policy, Commercial Terms of Service, and Privacy Policy (Anthropic, "Commercial Terms of Service - Effective May 14, 2024"; Anthropic, "How long do you store personal data?", published between the end of April 2024 and mid-June 2024, archived on 22 June 2024; Anthropic, "Usage Policy," archived on 13 May 2024). The three most clearly EU-relevant changes were (i) designating Anthropic Ireland, Limited as the data controller for users in the EEA, UK, and Switzerland, (ii) a 30-day default data retention period, and (iii) new disclosure requirements and more granular high-risk-use-case categories, though we cannot confirm from public sources whether these were new for the EU launch or pre-existing changes communicated alongside it. Confidence is low because while the incorporation of an Irish subsidiary and other changes are consistent with GDPR compliance work, Anthropic has not confirmed this was the reason for the delays and non-releases, and we cannot distinguish regulatory preparation from general commercial readiness for the European market by a then still relatively young organization.

[53] Meta AI, "BlenderBot 3: A 175B Parameter, Publicly Available Chatbot That Improves Its Skills and Safety Over Time," 5 August 2022.

[54] Meta, "What's New Across Our AI Experiences," 6 December 2023. The first international expansion did not occur until April 2024 (Meta, "Meet Your New Assistant: Meta AI, Built With Llama 3," 18 April 2024), suggesting the UK and EU exclusion at this stage reflects a staged product rollout.

The largest group of delays and non-releases concerns LLMs powering Meta's AI assistant across its digital platforms (WhatsApp, Messenger, Facebook, Instagram, and the Meta.AI web app). From the Llama 3 release in April 2024 onwards, the UK and EU were consistently excluded from these releases. Both the UK Information Commissioner's Office (ICO) and the Irish Data Protection Commission (DPC) had raised concerns about Meta's plans to train models on UK and EU user data, requesting that Meta pause its training plans in June 2024.[55]

The UK gained access to the text-only version of Llama 3.2 powering the Meta AI assistant 14 days after the US release. The short UK delay appears to reflect the fact that, in September 2024, the ICO publicly acknowledged changes Meta had made to its training plans – clearing the regulatory path for Meta to release in the UK.[56] The EU did not receive access until March 2025, when Meta released a version of Llama 3.2 explicitly not trained on EU first-party data.[57]

A related set of non-releases concerns Meta's open source multimodal Llama models. The Llama 3.2 Vision models and Llama 4 multimodal models were made available in the UK but explicitly excluded EU-based individuals and companies under Meta's Community License Agreement, while text-only models in the same families remained available globally.[58] Meta publicly cited the "unpredictable nature of the European regulatory environment" as the reason.[59] The UK was not excluded, possibly because the ICO had indicated it would not block Meta's training plans,[60] while the equivalent EU process remained unresolved.

There appears to be a distinct and more severe regulatory barrier associated with the image-input modality across all of Meta's products. Image-input Llamas for the AI assistant on Meta's digital platforms has never been released in the UK or EU. Image-input Llamas powering Meta Glasses' visual AI features faced delays of over 200 days in both jurisdictions. We speculate this may be due to the stricter GDPR processing requirements associated with image data. Photographs can reveal protected characteristics – such as racial origin, health status, or religious affiliation – without these being explicitly stated, potentially triggering special-category processing obligations under Article 9.[61] This is relevant to images in training data and images uploaded by users or captured in real time by Meta Glasses.

[55] ICO, "Statement in Response to Meta's Plans to Train Generative AI with User Data," 14 June 2024; DPC, "The DPC's Engagement with Meta on AI," 14 June 2024.
[56] ICO, "Statement in Response to Meta's Announcement on User Data to Train AI," 13 September 2024. The ICO explicitly stated it had "not provided regulatory approval."
[57] Anna Dack, Meta's innovation communications manager, quoted in *TechCrunch*, "Meta AI Is Finally Coming to the EU, but with Limitations," 20 March 2025.
[58] Meta's Llama 3.2 Community License Agreement explicitly excludes EU individuals and companies from using the multimodal (vision) models, while text-only models in the same family remain available (Meta, "Llama 3.2 Acceptable Use Policy", GitHub, archived on 27 September 2024).
[59] Fried, Ina, "Scoop: Meta Won't Offer Future Multimodal AI Models in EU," *Axios*, 17 July 2024.
[60] ICO, "Statement in Response to Meta's Announcement on User Data to Train AI," 13 September 2024.
[61] Under the Court of Justice of the European Union's inference doctrine, any processing that is "liable to indirectly reveal" protected characteristics falls within Article 9 regardless of the controller's intent (Case C-184/20, *OT v Vyriausioji tarnybinės etikos komisija*).

Google

Google also appears to have encountered barriers to releasing LLMs in the UK and EU due to data protection regulation – though much less significantly than Meta. Its delays and non-releases fall into two main categories: early chatbot launches where regulators directly intervened, and a series of free-tier API releases that were never released, due to legal barriers that made doing so less commercially desirable.

The clearest cases of regulatory intervention involve Google's early chatbot releases. In June 2023, the Irish DPC blocked the EU launch of Bard after Google failed to provide a data protection impact assessment.[62] Google delayed the EU launch by 64 days, after making a set of changes regarding transparency, user controls, and an opt-out for data collection.[63] The subsequent launch of Gemini 1.0 Pro in December 2023 was also delayed – by 13 days in the UK and 57 days in the EU – with Google's general manager for Bard stating that it was "working with local regulators," though the specific barriers were not disclosed.[64]

Most of Google's remaining non-releases relate to free-tier API releases of PaLM 2, Gemini 1.0, and Gemini 1.5 models during 2023 and 2024. We speculate that these were driven by data protection rules that prohibited Google from processing user inputs for model training – a key commercial incentive of providing free model access to developers. Rather than creating a legal barrier to release, these rules could have made the free tier commercially undesirable in the UK and EU. Google's eventual decision to update its terms of service in May 2024 – explicitly applying the paid-tier prohibition on training-on-user-inputs to users in the UK, EU, Switzerland, and other European Economic Area (EEA) countries on the free tier – appears to have enabled the release of Gemini 1.5 Pro Experimental with a 22-day delay in August 2024.[65] One possible interpretation is that Google initially withheld free-tier releases while waiting to see whether EU and UK data protection authorities would eventually permit the use of free-tier user inputs for model training. Once it became clear they would not, and as competitive pressure grew, Google accepted the less commercially attractive terms in order to release the model.

---

[62] Graham Doyle, Irish DPC Deputy Commissioner, quoted in *TechCrunch*, "Google Delays EU Launch of Its AI Chatbot After Privacy Regulator Raises Concerns," 13 June 2023.
[63] Lomas, Natasha, "Google's Bard and Other AI Chatbots Remain Under Privacy Watch in the EU," *TechCrunch*, 13 July 2023.
[64] Sissie Hsiao, Google VP and General Manager of Bard, quoted in *The Guardian*, "Google Says New AI Model Gemini Outperforms ChatGPT in Most Tests," 6 December 2023. See also Heikkilä, Melissa, and Heaven, Will Douglas, "Google DeepMind's New Gemini Model Looks Amazing – But Could Signal Peak AI Hype," *MIT Technology Review*, 6 December 2023.
[65] Google updated the Gemini API Additional Terms of Service to require that users in the European Economic Area, Switzerland, or the UK either use Paid Services or have the Paid Services data-handling terms applied to their usage on the free tier. The substantive effect was to prohibit Google from using these users' free-tier prompts and responses to train models or improve products – bringing the free-tier terms for these jurisdictions into line with the paid-tier terms, while free-tier users elsewhere remained subject to terms permitting training on user inputs. See Google, "Gemini API Additional Terms of Service," archived on 18 May 2024).

OpenAI

OpenAI has no examples of non-releases and its delays are generally short. Delays fall into two broad categories: staged international rollouts that appear unrelated to UK- or EU-specific regulation – including cases that plausibly reflect compute or capacity constraints – and a smaller number of delays where regulatory factors appear to play a role, though we find no strong evidence of a specific barrier in any case.

Two early delays appear to reflect phased global releases rather than jurisdiction-specific barriers. GPT-3.5 Turbo Alpha was delayed uniformly across all non-US countries, with the UK and EU gaining access at the same time as many other jurisdictions. Operator, powered by the GPT-4o Computer Use Agent model, followed a similar pattern for the UK, which launched alongside other non-EU markets as part of a broader international expansion. This is consistent with caution around the risks of agentic capabilities, or with compute constraints, rather than UK-specific regulatory barriers.

A larger group of delays shows a pattern more suggestive of regulatory factors, with the EU and UK facing the same delay (GPT-4 image input), the EU facing longer delays (GPT-4o audio; GPT-4o Computer Use Agent), or the EU facing delays but the UK gaining immediate access (GPT-4o real-time audio and video; ChatGPT agent). The exclusion of European countries specifically, rather than a broader set of jurisdictions, suggests that common European data protection laws may be a factor. The cases also cluster around new input modalities (image, voice) and novel capabilities (agentic browsing), which plausibly trigger additional compliance review under GDPR – though our confidence is low, as OpenAI has not elaborated on the specific barriers in any case.

One case – o3 Deep Research, delayed 3 days in both the UK and EU – is most plausibly explained by compute constraints. OpenAI described the feature as “very compute intensive”, and a short delay that was identical in the UK and EU fits a capacity-limited rollout better than a regulatory barrier, which in OpenAI's record tends to be longer and/or EU-specific.

Anthropic

Anthropic has the fewest delays and non-releases of the four companies in the dataset and is the only company with no delays or non-releases in the UK. All six cases occurred in the EU during 2023 and 2024 and all are assessed with low confidence.

Anthropic did not release Claude 2 or Claude 2.1 to the EU and delayed release of the Claude 3 family (Haiku, Sonnet, and Opus) by approximately two months. We speculate that EU data protection requirements – particularly the need to establish a legal entity, designate a data controller, and create GDPR-compliant consumer terms – made the EU harder to enter relative to other markets for a then relatively young organization, leading to deprioritization

rather than an explicit block. When Anthropic did enter the EU, it incorporated Anthropic Ireland Ltd as its EU data controller and launched the Claude 3 web app in Europe with updated consumer terms and a 30-day default data retention period.[66]

Notably, the Claude API publicly launched simultaneously in the EU and US in March 2024,[67] suggesting there was no hard legal barrier to deploying Claude 2, 2.1, and 3 models themselves; the web app delay appears to reflect the time needed to establish consumer-facing compliance infrastructure.

No regulator is known to have intervened at any point, and Anthropic has not publicly stated the reasons for the delays and non-releases. It is therefore difficult to distinguish regulatory preparation from general commercial readiness for the European market.

[66] Anthropic, "Claude is now available in Europe," 14 May 2024; Anthropic, "Commercial Terms of Service - Effective May 14, 2024"; Anthropic, "How long do you store personal data?", published between the end of April 2024 and mid-June 2024, archived on 22 June 2024.
[67] Anthropic, "Introducing the Next Generation of Claude," 4 March 2024.

Methodology

We collated a dataset of 375 large language model[68] public releases between 06/2018 – 05/2026 from OpenAI, Anthropic, Google DeepMind, and Meta.[69] The dataset can be accessed here. Please get in touch with the authors if you are interested in developing this dataset.[70]

For each model release, we compare the release date in the US with those in the EU and UK[71] using public sources.[72] All the releases up until 31 October 2025 were manually identified and verified by human reviewers; releases from November 2025 onwards were initially identified by Claude Cowork using our established scoping criteria and methodology, and were then subject to several rounds of human review.

A "model release" can include different releases of the same model. This includes the mode of release (e.g. via API, web/ mobile app, open source, or integration into other products), release stage (e.g. general release, public preview, alpha), additional modalities (e.g. vision, computer use), context length, and new versions of the model.

For example, the dataset includes four model releases for Claude 3.5 Haiku, based on variations in the mode of release and input modalities:

Model	US Release Date	Mode of Release	Modalities (input)
Claude 3.5 Haiku	04/11/2024	API	Text
Claude 3.5 Haiku	12/12/2024	Consumer App	Text
Claude 3.5 Haiku	24/02/2025	API	Text and Images
Claude 3.5 Haiku	24/02/2025	Consumer App	Text and Images

For each example of a model that was delayed or not released to the UK or EU, we used public sources to determine the most likely cause of the delay or non-release. In most cases, we were only able to assign low or medium confidence to our assessment.

[68] We exclude models that do not generate text at all such as dedicated image, audio, or video generation models (e.g. OpenAI's Sora) – this was purely to keep the scope of the analysis manageable.
[69] These companies have released the most significant models, as judged by the Epoch AI Capabilities Index. Intended to be indicative rather than comprehensive. Future studies could broaden the scope to other companies. The cut-off date of data collection is 31 May 2026. We did not define any specific beginning date for the data collection: 2018 was the year when one of the four companies first released an LLM publicly. Private releases, i.e. invite-only or accessible only via a waiting list, are not part of our dataset, because it is impossible to verify whether all regions were granted access within the same delays after joining a waitlist.
[70] Please write to contact@governance.ai
[71] Generally EU countries have the same release dates; variation between countries is recorded in the dataset and within footnotes in this post.
[72] This includes company blogs, social media statements and time-captured FAQ pages, YouTube videos, and forum discussions.

We do not consider something a delay and non-release where a model was not available for free, but was available via a paid subscription; or where a model was not available via mobile app, but was available via a web app. In both cases, we interpret this as UK/EU users having sufficiently similar access to the model as US users to consider it a simultaneous release.

We analyze public sources to interpret the most likely reason for the delay or non-release. Where the information is readily available, we analyze which countries, beyond the US, were given access and in what sequence – which can imply the presence or absence of jurisdiction-specific factors. For example, if it is only the EU that is left out of a global rollout, then we interpret this as a strong signal of EU-specific factors. Likewise, if the EU and UK are part of a large group of countries that encounter the same delay, this is a strong signal of the absence of jurisdiction-specific factors. However, a systematic tracking and analysis of LLM releases in other non-UK and non-EU countries was out of scope of this report. Whenever we mention that only European countries are left out of a global rollout, we do not take into consideration the countries that have usually been excluded from US releases for geopolitical reasons during 2018-2026, notably China, North Korea, Belarus, Russia, Cuba, Iran, and Syria.

Companies were rarely explicit about the reason behind a delay or non-release – and even in these instances it was usually at a high level – e.g. "due to the complex regulatory environment" or "we are working with regulators" – rather than being specific about a particular regulatory barrier. We also assume that public claims about regulatory factors should not be wholly relied on, as companies may have incentives to overstate the impact of regulatory barriers as a means of pressuring regulators and policymakers to soften regulation.

We found even less direct evidence about non-regulatory factors, presumably because companies would be less incentivized to reference these publicly.

In some cases there is strong evidence of a specific regulatory barrier, because a regulator has explicitly referenced it in engaging with the company, e.g. the Irish Data Protection Commissioner's engagement with Meta concerning its use of Facebook and Instagram data for model training, or direct prevention of Google's release of Bard due to insufficient data protection documentation.

Mostly, though, we rely on inference about which particular regulatory barriers might be at play given the nature of the model. For example, we see a strong pattern in Meta's delays and non-releases relating to models with an image input modality and thus we try to infer how such a modality might pose distinct GDPR compliance challenges.

Limitations and Future Work

Though we find fairly modest evidence of delays and non-releases here, this may change as new regulations are put in place and start being enforced, notably the General-Purpose AI provisions of the EU's AI Act. We would welcome other researchers to reach out to us if they are interested in continuing the work of populating this dataset as new models are released and the regulatory environment changes.

Our analysis may also have underestimated regulatory barriers by focusing on frontier AI companies, which, among other things, have more developed legal teams. Future research could look at a broader set of companies, including startups, non-US AI companies, and companies that build applications and products using AI models, rather than training those models themselves (e.g. Apple).

Further, our analysis has focused primarily on LLMs, which we have defined as models that can (at least) produce text outputs. The pattern may look significantly different for image-, audio-, and video-generation models. We would welcome other researchers to expand the dataset to include such models.

We have focused on the EU and UK, as two important markets for the release of frontier models. The same analysis could be useful for other jurisdictions such as Australia, Canada, the US (especially looking for state-level differences), Brazil, and Japan. Additionally, our analysis can be further expanded by investigating whether the UK and EU overseas territories have experienced the same delays and non-releases where the data is available.

The weakest part of our analysis is our assessment of the causes of non-releases and delays. Policymakers may be well-positioned to dig deeper into those causes, and researchers with a deeper understanding of the relevant regulatory instruments may be able to identify and correct our inevitable mistakes.

About the Authors

John Lidiard
Research Scholar, GovAI

John is a Research Scholar focusing on UK AI policy. He worked in the UK government as a policy advisor on digital technology regulation and helped to set up the AI Security Institute. He has a BA in English from Oxford, and MA in International Security from King's College London.

Oleksandra Vereschak
Research Associate, Future Impact Group

Oleksandra Vereschak, PhD, is a researcher specializing in the ethics and governance of AI. She completed this work as a Future Impact Group Research Associate. Her research covers human perception of AI and the impact of AI on users and model subjects. She holds a PhD in Human-AI Interaction from Sorbonne Université.

Tom Gibbs
Research Assistant, GovAI

Tom is a Data Scientist at British Airways and an AI Red Teamer at Trajectory Labs, where he develops adversarial evaluations and prompt injection attacks for frontier labs. His published research covers multi-turn and compositional jailbreak methods. He is a Research Assistant at GovAI and holds a BSc in Economics along with an MSc in Machine Learning.

Markus Anderljung
Director of Policy and Research, GovAI

Markus leads research at GovAI. He served as one of the Vice-Chairs drafting the EU's Code of Practice for General Purpose AI, and was previously seconded to the UK Cabinet Office as a Senior AI Policy Specialist. He is also an Adjunct Fellow at the Center for a New American Security and a member of the OECD AI Policy Observatory's Expert Group on AI Futures.

Acknowledgements

We are grateful for valuable conversations and feedback from the following individuals: Alan Chan, Noemi Dreksler, Elias Groll, Aidan Homewood, Zaheed Kara, Cornelia Kutterer, Sam Manning, José Luis León Medina, Peter Wills, and the Future Impact Group.

About GovAI

GovAI is a 501(c)(3) non-profit organization. Our mission is to help decision-makers navigate the transition to a world with advanced AI by producing rigorous research and fostering talent. Researchers at GovAI work on a wide range of topics, with a particular emphasis on the security implications of frontier AI.