
Integration of Generative AI in the Digital Markets Act: Contestability and Fairness from a Cross-Disciplinary Perspective

Ayse Gizem Yasar
London School of Economics

Andrew Chong
University of California

Evan Dong
Cornell University

Thomas Krendl Gilbert
The New York Academy of Sciences

Sarah Hladikova
Tufts University

Carlos Mougán
University of Southampton
The Alan Turing Institute

Xudong Shen
National University of Singapore

Shubham Singh
University of Illinois Chicago

Ana-Andreea Stoica
Max Planck Institute for Intelligent Systems, Tübingen

Savannah Thais
Columbia University

Abstract

This paper explores whether generative AI should be explicitly included as a core platform service under the EU’s Digital Markets Act (DMA). We argue that it should be, as generative AI is becoming a significant gateway between businesses and users, similar to other digital platforms, and current regulations may not adequately address the specific challenges it presents. The paper is divided into four parts: (i) *Platformization of Generative AI*: Generative AI is emerging as a new type of digital platform, either as infrastructure or as an augmentation to existing services. (ii) *Gatekeeping Characteristics*: Some companies developing generative AI models exhibit gatekeeping traits like first-mover advantage, capital and data resources, and system integration, leading to concerns about fairness and contestability in the market. (iii) *Generative AI as a Core Platform Service*: We suggest amending the DMA to include generative AI as a core platform service, aligned with the AI Act’s definition of general-purpose AI systems. (iv) *Obligations for Generative AI Gatekeepers*: Current DMA obligations may not fully address the fairness and contestability issues in generative AI. We suggest that new obligations may be needed for generative AI gatekeepers under the DMA.

1 Introduction

The European Union’s response to platform size and power in digital markets goes beyond traditional competition law enforcement. It is only when market power is abused that EU competition law intervenes. However, there is an increasing concern that ‘a few large platforms act as gatekeepers between businesses and users and enjoy an entrenched and durable position, often as a result of the creation of conglomerate ecosystems around the core platform services, which reinforce existing entry barriers.’ The underlying logic is that these entrenched positions lead to unfair behaviour vis-à-vis business users of these platforms, as well as reduced innovation and contestability in core platform services.

These concerns led to the creation of a new European regulation that goes beyond traditional competition law rules: the Digital Markets Act (DMA). It aims to keep the entry barriers low in digital markets and ensure fair game between ‘gatekeepers’ and their smaller rivals that depend on gatekeepers’ services. The regulation covers some well-established platform services like operating systems, messaging platforms, and online advertising. The DMA is silent on AI and in particular generative AI, but there is an increasing concern that gatekeepers are emerging in generative AI applications.

In this paper, we argue that generative AI should be included in the DMA as a core platform service. Our argument relies on an analysis of the gatekeeping potential of companies that develop and deploy generative AI models. We argue that generative AI services increasingly act as core platform services within the meaning of the DMA, yet cannot be categorised as one of the current examples of core platform services and must therefore obtain a category of their own. We illustrate the ways in which certain generative AI companies may undermine fairness and contestability as defined in the DMA. We draw on the landscape of existing generative AI research, as well as characteristics of digital markets and past business practices by core platform service providers that led to the adoption of the DMA.

The paper is structured as follows: In Section 2, we survey current generative AI models and describe the pathways through which they become *gateways* between businesses and users within the meaning of the DMA. More specifically, we identify how certain types of generative AI models emerge as *novel types of digital platforms*, fundamentally different from existing platforms. We argue in Section 3 that companies that develop large generative AI models strongly exhibit four main characteristics of digital services which facilitate the emergence of gatekeepers: first-mover advantage, capital-intensive resources, data resources, and system integration. In Section 4, we outline a regulatory proposal, arguing that the DMA has the potential to address contestability and fairness issues in the generative AI industry if it is amended to include generative AI as a core platform service. Finally, we provide examples of existing gatekeeping obligations that would apply to generative AI if it were directly included in the list of core platform services. Legal foundations of our proposal are provided in Appendix A.

2 Platformization of Generative AI

The DMA applies to ‘core platform services’ that act as an important gateway for business users to reach end users. It provides a closed list of core platform services in Article 2(2) (see Appendix A). This section sets out the different ways in which generative AI applications act as platforms within the meaning of the DMA.

Generative AI has enabled novel digital applications and significantly enhanced existing ones. The technologies underlying generative AI systems and generative AI applications are increasingly used as platform infrastructure by business users. As a result, certain generative AI technologies and applications are starting to act as gateways between business users and end users within the meaning of the DMA. Below, we first discuss the transformation of generative AI into platform infrastructure, followed by use cases where generative AI augments existing services.

2.1 Digital Platform Infrastructure

A few companies in the generative AI industry are developing the critical infrastructure for generative AI-enabled services. Central to this infrastructure is what is known as a foundation model. These are large-scale AI models that have been pre-trained on diverse data at scale and are designed to be versatile (Bommasani et al., 2021). Analogous to how cloud computing provides ‘infrastructure to support and enable functionality in services offered by others’ (European Commission, 2020a), foundation models function as the infrastructure for a wide range of generative AI applications and services because they provide the fundamental capabilities of generative AI.

Within the foundation models field, a distinction is often made between *frontier* models and *behind-the-frontier* models. Frontier models, as defined by the Frontier Model Forum,¹ is ‘large-scale machine-learning models that exceed the capabilities currently present in the most advanced existing models’ (Microsoft, 2023b). On the other hand, models that do not reach this level of capability are considered behind-the-frontier models. This distinction is crucial because frontier models can suddenly obtain capabilities that behind-the-frontier models lack, a phenomenon known as

¹The Frontier Model Forum was established by OpenAI, Anthropic, Google, and Microsoft. See <https://openai.com/blog/frontier-model-forum>

emergent abilities (Wei et al., 2022; Schaeffer et al., 2024). Such emergent abilities give frontier models a considerable competitive advantage, positioning them as essential infrastructure that offers functionalities unattainable by other models.

Key to this process is the adaptability of foundation models, meaning that developers of generative AI applications need not develop their own models. Instead, a pre-trained foundation model can be built ‘on top of’ to develop various generative AI applications. For example, OpenAI’s Generative pre-trained transformer (GPT) models, besides underpinning OpenAI’s own chatbot product ChatGPT, have also been customized by startups like Harvey to assist lawyers in accessing and generating legal content (Grady and Curnin, 2023), and by Jasper to aid in marketing content creation (Jasper, 2023). This paradigm has fostered a delineation between the platform layer—which provides foundational models—and the application layer, which focuses on building user-facing products. Consequently, the foundation model emerges as a novel type of digital platform that provides essential generative AI functionalities for business users to deliver generative AI applications to end users.

Establishing foundation models as platforms is becoming a strategic imperative across the generative AI industry, especially for the developers of frontier models. For example, OpenAI provides (1) API access (OpenAI, 2020) to proprietary GPT models, (2) GPTs, which enable lightweight customization of ChatGPT for specific purposes (further discussed in the next sub-section), and (3) a developer product called Foundry (Wiggers, 2023), which allows for the fine-tuning of GPT models and static allocation of compute resources to support large-scale inference. In a similar vein, both Amazon Bedrock (Amazon, 2023) and Google Vertex AI (Google, 2023) offer a broad collection of foundation models along with the tools required to develop generative AI applications.

2.2 Augmented Services

The capabilities of generative AI come alive when interacting with users. This is possible through a new breed of user interfaces such as Google’s Gemini, which integrates inert model capabilities into the behavioural affordances of a system. As these capabilities become dynamic and multimodal, combining various tasks at once, the user experience in digital services is transformed. This qualitatively raises the stakes for legacy providers across various service domains such as search engines.

The significance of chatbots lies not just in their technical proficiency at simulating human linguistic patterns, but in the material threat they pose to market actors who offer the services they augment. This is most clear in the panic ChatGPT caused to Google and Microsoft, which had been developing their own generative AI models internally for years but had not released them publicly. OpenAI’s success was in developing a user interface (ChatGPT) that, when paired with a large language model, became a compelling alternative to the web search interfaces that have defined the internet experience over the past quarter century. Google acted quickly, merging DeepMind and Google Brain in April 2023 to become Google DeepMind, while Microsoft immediately made a multibillion-dollar investment in OpenAI itself and integrated ChatGPT’s underlying language model into its Bing web search tool.

There are signs that generative AI chatbots and, more broadly, generative AI applications are turning into gateways for business users to reach end users. A significant number of institutions have started to employ generative AI models in servicing users for a variety of applications, including in the European Union (Macaulay, 2023). Generative AI has already become commonplace as a tool used by businesses to serve customers in both small businesses (Soni, 2023) and large businesses (Kanbach et al., 2023). For example, Adobe Sensei use genAI integration in creative software tools (Insights, 2023; Adobe, 2023), whereas Siemens integrates genAI in medical diagnoses services (Siemens, 2023).

In keeping with this trend, OpenAI has enhanced ChatGPT by introducing third-party plugins. These plugins are typically developed by businesses that specialize in specific digital services. For instance, Expedia and KAYAK have created plugins to integrate with their own websites, enabling ChatGPT users to search for flights and accommodations. While these plugins extend the capabilities of ChatGPT, they also strengthen its role in directing and managing traffic for business users, making ChatGPT more platform-like. OpenAI has also introduced another product called GPTs, which allows developers to customize ChatGPT for specific purposes, coupled with a GPT store for distributing custom GPTs. Against this background, OpenAI’s gateway role is twofold: The original ChatGPT serves as the generative AI infrastructure, akin to how cloud computing provides the necessary infrastructure to support and enable the functionality of various services. The introduction of the GPT Store marks OpenAI’s strategic shift towards a platform-centric ecosystem. While the monetization pathways are still being defined, OpenAI’s CEO, Sam Altman, has indicated ‘[w]e’re going to pay

people who make the most used and most useful GPTs [i.e., customized ChatGPTs for specific purposes] with a portion of our revenue’ during OpenAI’s first Developers Day (Coldevey, 2023).

Chatbots raise the question of whether they could be categorised as virtual assistants, which are a type of core platform service listed in the DMA. Although there might be a conceptual resemblance, chatbots diverge from virtual assistants in important ways. According to Art. 2(12) DMA, a virtual assistant is ‘a software that can process demands, tasks or questions, including those based on audio, visual, written input, gestures or motions, and that, based on those demands, tasks or questions, provides access to other services or controls connected physical devices’. While chatbots do process user prompts in terms of ‘demands, tasks or questions’, their primary function is neither to ‘provide access to other services’ nor to ‘control connected physical devices’, despite their capacity to perform these tasks. Fundamentally, what generative AI chatbots provide is *the generated responses (or, more broadly, data) themselves*, rather than access to other distinct services. The DMA defines virtual assistants in this way because it is motivated by virtual assistants in the IoT industry, such as Amazon’s Alexa and Google’s Assistant. The drafters of the DMA did not foresee the emergence of generative AI-based chatbots. There is some room to argue that to the extent a generative AI chatbot is designed to refer users to other services, it falls under the virtual assistant definition. However, to the extent generative AI applications serve as infrastructure to build other generative AI applications, they fall outside the scope of ‘virtual assistants’.

3 Generative AI and Gatekeeper Characteristics

The generative AI sector has recently experienced significant growth, with both established tech companies and myriad startups entering the market to develop and/or deploy foundation models and new applications in generative AI. Some of these companies hold a significant competitive advantage due to their extensive data resources, specialized hardware architectures, vertical integration, previous talent acquisition, financial clout, and accumulated knowledge through prior product releases. Many of these advantages—such as data resources, financial clout, integration, and the importance of early movers—have been present in digital markets in general and given rise to the gatekeeping positions that the DMA now seeks to address.

In this section, we catalogue four major advantages exhibited by several generative AI companies: (i) first-mover advantage, (ii) computing resources, (iii) data collection, and (iv) multimodal system integration. We explain these dimensions where appropriate by drawing analogies to pre-generative AI instances of gatekeeping. Without regulatory intervention, these new advantages will likely solidify strong positions within the generative AI services landscape, mirroring patterns observed in other digital markets.

3.1 First-Mover Advantage from Novel Product Features

Some generative AI companies, especially those developing large language models, possess a first-mover advantage. First-mover advantage is defined as a competitive advantage gained by being the first one to introduce a product or service in a market, allowing it to build brand recognition and customer loyalty (Lieberman and Montgomery, 1988). A preliminary distinction must be drawn between being the *first* vs. being the *first-mover*. The first company or team to develop a model with new capabilities is not necessarily the first to capitalize on those capabilities. While a first-mover advantage is not problematic as such, in markets exhibiting strong network effects, it can lead to entrenched and durable positions.²

The company exhibiting a clear first-mover advantage in the generative AI industry is OpenAI. Google developed the transformers architecture (Vaswani et al., 2017) that became essential in large language models, yet this did not cement Google’s position as a first-mover within generative AI. Although Google developed and open-sourced several large language models that used the previously developed transformers architecture, such as BERT (Devlin et al., 2019), LaMDA (Thoppilan et al., 2022), and PaLM (Chowdhery et al., 2023; Anil et al., 2023), neither model took off as a large-scale product used by the public. There are two reasons for this: first, BERT (the first large language model developed by Google) was designed and developed primarily as a scientific tool to demonstrate the

²See paragraph 79 in the Impact Assessment European Commission (2020a).

possibility of tokenization and not as a large engineering project (Vipra and Korinek, 2023). Second, neither BERT nor LaMDA had a public, low-friction user interface.³

OpenAI created GPT-2 in 2019 (Radford et al., 2019), shortly after the release of BERT, with an easy-to-use interface but limited to a subset of users that requested access via their own website. The stated intent of the limited access policy was to prevent ‘malicious’ use from spreading in the name of due diligence (OpenAI, 2019), the collateral effect of this policy was to induce scarcity around a new kind of resource: the unrivalled capabilities of large language models. GPT-2 was comparable to other state-of-the-art models and did not constitute a leap in performance. However, the limited access policy created a new dynamic in which the research community went from ‘open-by-default’ to scarcity around research access to large language model capabilities (Vincent, 2019).⁴ Because OpenAI put the emphasis on safety concerns emanating from the state-of-the-art capabilities of the model, the limited access policy also created an early standard for understanding large language models’ capabilities and supposed risks.

GPT-3 introduced a new user interface, ChatGPT, alongside a significant increase in conversational capabilities. Notably, there is a distinction with Anthropic’s Claude, which exhibited similar capabilities to GPT-3 but was tested internally within the company’s Slack channels rather than made available to the public. GPT-4 further extended these capabilities by introducing new features, such as larger output sizes and system integration with Bing’s browser services. Additionally, it introduced the ChatGPT Plus service, offering priority access to new features for a monthly subscription fee of \$20/month.

OpenAI continues to release novel GPT models, and each release represents an upgrade in their capabilities and a strategic evolution in OpenAI’s place in the market. OpenAI carefully plans each model’s debut, introducing new features that are initially scarce before becoming widely available and, ultimately, monetized. This strategy highlights OpenAI’s efforts to enhance its market influence with each new model. Specifically, GPT-2 laid the groundwork by creating a demand for research access, effectively setting a new standard for accessing and understanding these models in both academic and commercial settings. Following this, GPT-3 focused on capturing a larger share of computing resources, while GPT-4 advanced into monetization, linking directly to the concept of controlling access as outlined by the Digital Markets Act (DMA). OpenAI is actively pursuing regulatory measures to maintain its competitive edge by safeguarding its technology from open-source distribution and reinforcing its market position.

Moreover, securing a first-mover position often means defining the benchmarks for software performance and ethical standards within the industry. For instance, the deployment and assessments of GPT-4 have established the criteria for evaluating the effectiveness of large language models. Such standards are now the baseline for comparing new methodologies, as seen in the widespread reference to GPT-4 in recent studies (Achiam et al., 2023; Saparov et al., 2024; Gemini Team et al., 2023; Li et al., 2024). These comparisons extend to performance on standardized tests, human evaluations, and also the significant expansion of the user community (Duarte, 2024), illustrating GPT-4’s role in setting the industry’s performance and value expectations.

3.2 Capital-Intensive Resources Required for Performance Optimization

State-of-the-art models in generative AI require a significant amount of computing power, as measured through the number of floating point operations or FLOPs.⁵ Computing power is needed to process the training data through a model: on the one hand, the model undergoes updates through its weights and parameters; on the other hand, the data that is used for training is being read, processed, and re-used in many iterations of the model updates. Thus, the larger a model is (i.e. the more parameters it has) and the larger the training data is, the more computing power is needed for

³While BERT has garnered academic interest, it is often criticized for its suboptimal tokenization methods.

⁴As a subset of the coauthors can attest, although elements of BERT were language-dependent (it was initially English-only), its architecture was made open-source and was understood by the research community as comprising an ongoing scientific and technical agenda being carried out by relevant teams at Google. As such, ongoing access was informally maintained by Google and, in effect, guaranteed by having a standing research affiliation with one of its labs. For these reasons, despite being led by a corporate entity, research on language models did not face formal organizational barriers before GPT-2. For more information, see Devlin and Chang (2018) and references therein.

⁵Not to be confused with FLOP/s, which is the measure for computing performance defined as the number of floating point operations per second.

successful training. As AI models have undergone significant growth in model size, complexity, and training data size, generative AI is at the forefront of computing power usage.

Computing power is critical for two tasks in generative AI: training and inference. Training is used to update the model parameter and weights through several iterations by minimizing an objective function, whereas inference is defined as the task where the model outputs a response based on a query (e.g. a prompt given to ChatGPT outputs a response, which is the inference step). Training uses computing power in terms of both required memory storage and FLOPs. For example, as a recent report from Andreessen-Horowitz argues (Appenzeller et al., 2023), the memory requirement during training for a model like GPT-3 that has 175 billion parameters is over a terabyte. State-of-the-art hardware options, such as NVIDIA’s A100s, currently offer 80GB of high bandwidth memory, meaning that multiple units are needed for training models of such size (NVIDIA, 2021). For example, ByteDance is training LLMs with GPU clusters at the scale of more than 10,000 GPUs (Jiang et al., 2024). Similarly, the amount of computing power needed for training (FLOPs available) is directly linked to the training time (e.g. a 175 billion parameter model would take decades to train on a single GPU). In contrast, the computing power needed for inference is significantly less (an inference prompt is tiny, incomparable in size to the training data, and would take less than 1 second on an A100 machine).

State-of-the-art models have seen a massive increase in the number of parameters used (Achiam et al., 2023; Ramesh et al., 2022). Furthermore, the size of the model is not the only variable in determining the cost of development, as all models in generative AI (larger or smaller) use very large datasets for training. More specifically, the larger the training dataset, the more epochs a model needs to undergo in the training phase, thus increasing the number of FLOPs used. Thus, it is the training data size and the model size together that create a huge demand for computing power, and the latest hardware specifications are needed to accommodate these evolving computational requirements.

However, acquiring and using the necessary GPUs for training large models is no easy task for two reasons: (1) increasing scarcity in their availability (e.g. Nvidia’s A100s have been in short supply in 2023 (Appenzeller et al., 2023)) and (2) large training costs. For example, estimates state that using cloud providers, the cost of training GPT-3 is of the order of \$4 million (without the added inference costs) (Thompson et al., 2021), while the cost of training GPT-4 is close to \$100 million (Knight, 2023). Of course, the exact costs are not publicly available and may be slightly reduced due to the use of in-house infrastructure or optimization methods. On this last point, even with an in-house infrastructure, training GPT-4 is estimated to have needed around 25,000 A100 GPUs (for reference, the latest 8-GPU system from Nvidia, NVIDIA DGX A100 has an estimated cost of \$200,000 in 2023). Smaller generative AI models, such as Stable Diffusion, have a reported cost of \$600,000 for training, according to the CEO’s public statements.⁶ These costs do not include the acquiring, storage, and maintenance costs regarding the training dataset, as well as inference costs, personnel costs, etc. It was the need for financial resources to secure the access to computing power that led OpenAI to change their business model and become a capped for-profit company.⁷

It is doubtful that the availability of small models and the decreasing costs of fine tuning are enough to counteract the increasing power of the lucky few who can afford and access the computing power necessary to train large generative AI models. The amount of compute needed in developing AI models has grown not only for large models, as regular-scale models have been doubling their compute usage in just 5.7 months while large models have been doubling their compute usage in 9.9 months (Sevilla et al., 2022). Fundamentally, this represents a change from the ‘Deep Learning Era’, as argued by Sevilla et al. (2022): on average, large-scale models use two orders of magnitude more computing power than classical deep learning models (see Figs. 1–3 in Sevilla et al. (2022)).

Meanwhile, there is evidence that smaller models cannot compete with the capabilities of state-of-the-art large generative AI models (Gudibandé et al., 2023). Despite the decreasing cost of fine-tuning, achieving state-of-the-art performance still requires a high budget, creating an entry barrier for potential players. Coupled with the scarcity concerns emanating from the global supply chain issues for manufacturing and shipping state-of-the-art GPUs, we are at a point where the sheer amount of computing power needed appears to favour the actors with immense financial and reputational resources in acquiring such GPUs. This entry barrier disproportionately affects smaller companies, public institutions, and universities, who often lack the financial resources to establish independent

⁶<https://twitter.com/EMostaque>

⁷Sam Altman described this shift in an interview with Lex Friedman, available at https://www.youtube.com/watch?v=L_Guz73e6fw. See also (Brockman et al., 2019).

large-scale generative AI systems (Vipra and Myers West, 2023; Murgia, 2023). Consequently, the concentration of cutting-edge computing power and expertise in the hands of a few players risks limiting player diversity and stifling innovation within the generative AI industry.

3.3 Data Resources

High-quality training and evaluation data are key for any genAI. While many datasets are publicly available or sold commercially, some generative AI companies gain an advantage by using private, customized datasets, which often include personal data. First, established platform service providers hold massive troves of data from *legacy applications* that can be used to train or evaluate generative AI models. For example, Microsoft owns GitHub, and a current lawsuit alleges that Microsoft used private GitHub repository data to evaluate its Copilot generative AI coding tool (Firm and Butterick, 2023). Additionally, as argued by FTC (2023), established players may benefit from ‘honed proprietary data collection tools and technologies for acquiring or scraping data’.

Usage data generated from interactions with AI applications (through prompts and other inputs) is also valuable. The more generative AI models are used and interacted with, the more their quirks or issues are tested, understood, optimized, and adapted. Several generative AI providers are in a position to collect usage data from both individuals and business users. While, for example, OpenAI claims to no longer use enterprise or API customer data to train its models (OpenAI, 2024a), it has in the past, possibly up until March 2023. Nor does this preclude the use of data from individuals. Although individuals in the EU retain the right to opt out of this data processing under the GDPR, the process to opt out is not user-friendly (OpenAI, 2024b,c), and it is unclear how many users do so in practice. Other large firms also use input and output data for model training. For example, Anthropic has an opt-in mechanism to use input data for model training (AnthropicAI, 2024).

Even when large firms do not directly use business user input and output data to train models (e.g. because the user has opted out), they still gain access to a vast trove of information and metadata about *how* business users use generative AI systems. For example, OpenAI reserves the right to run business data through content review systems and classifiers (OpenAI, 2024a). Due to the rapidly evolving nature of generative AI models and the large role played by RLHF and red teaming in improving generative AI model performance, this type of metadata may be quite useful in establishing entrenched and durable positions. For example, companies may be in a position to leverage the quality and scale of metadata to address concrete challenges in RLHF fine-tuning Casper et al. (2023), unlocking new technical capabilities unavailable to smaller competitors.

Another side-effect of such rich and complex data collection systems is the threat to end-user privacy. A recent data breach by Microsoft researchers exposed 38 terabytes of user data to the public (Ben-Sasson and Greenberg, 2023). Furthermore, OpenAI’s data collection practices have led to the ban of ChatGPT in March 2023, as it could not reliably prove that it did not collect personal information from children below 13 years of age (Weatherbed, 2023). Later in the year, a privacy researcher, Lukasz Olejnik, filed a complaint before the Polish data protection authority, alleging that OpenAI failed to produce details on how it collects personal data (Lomas, 2023). Since data is the primary raw ingredient in the recipe to create LLMs, generative AI providers might be incentivised to keep their data sources and collection practices out of the public eye.⁸ Privacy is not a direct objective of the DMA, and issues of opacity and security are already addressed in the GDPR and now also the AI Act. However, the DMA clearly emphasises the intersection of data collection practices, privacy and contestability, by explicitly mentioning the obligation of the gatekeeper to ‘protect privacy’ and encouraging ‘an adequate level of transparency of profiling practices employed by gatekeepers’.⁹

3.4 Integrated Systems

A growing trend in generative AI systems are integrated services, especially by companies already designated as gatekeepers under the DMA. A variety of services have started to integrate LLMs, including search engines (Microsoft, 2023; Reid, 2023), personal assistants, note-taking (e.g. Notion), and editing. Platform tendency for integration and the associated loss of user autonomy is not new: Google, Microsoft and Apple have all pushed for integrated systems in their now-established services,

⁸Restrictions on data-access or placing data behind paywalls exist KeyserSosa (2023); Nylén and Bass (2023)

⁹Recitals 35 and 72

such as search and email interfaces. As integrated systems that use generative AI become ubiquitous, their convenience and potential for creative endeavours trade off with user autonomy.

This integration goes both ways between genAI and other services. Both Google Search and Microsoft's Bing search engine have built functionalities drawing upon generative AI models—indeed, 'GPT-4 is already used in many places on Bing, including Copilot, Image Creator from [Microsoft] Designer, and already in regular web result ranking', according to an official Bing blog post announcing yet another GPT-powered integration, Deep Search ([Microsoft, 2023a](#)). Microsoft's Services Agreement confirms that as part of its AI services, which would include Bing powered by GPT-4,¹⁰ 'Microsoft will process and store your inputs to the service as well as output from the service, for purposes of monitoring for and preventing abusive or harmful uses or outputs of the service' ([Microsoft, 2023c](#)). In a similar vein, Google connected its Bard generative AI model to a suite of services such as Gmail, Google Docs, Google Maps, and YouTube with its Bard Extensions update in September 2023 ([Pinsky, 2023](#)). From the other end, ChatGPT gained 'browse with Bing' functionality in September 2023 ([Arunasalam, 2023](#)). This integration trend further muddies the distinctions between different services when it comes to data processing practices. As services and data become increasingly intertwined, it becomes difficult for a user to switch between platforms and transport their data and projects. As emphasised in the DMA (Recital 3), control over whole ecosystems allows gatekeepers to leverage their advantages from one area of activity to the other, which in turn reduces contestability.

4 Proposal for a New Core Platform Service: General Purpose AI Systems

The potential harms of AI are now materialising: the current race toward generative AI might lead to the emergence of a few large players at the expense of contestability and fairness in the generative AI industry. The DMA does not directly address this problem. Certain cases of generative AI applications might already be indirectly caught by the DMA when current gatekeepers integrate proprietary or third-party generative AI into their core platform services. Their standalone generative AI offering might also be caught. For example, the DMA specifically prohibits certain instances of personal data and service integration across different core platform services and other services provided by gatekeepers (e.g., Art 5(2)(b) and 5(8)).

However, without the explicit inclusion of generative AI as a form of core platform service, any protection provided by the DMA can only have a limited effect. Such indirect application does not address generative AI's specific contestability and fairness concerns. For one, unless generative AI is listed as a separate core platform service, any generative AI behemoths that do not themselves provide any other core platform service cannot be designated as gatekeepers under the DMA. Furthermore, the DMA's gatekeeper obligations will not even indirectly apply to existing gatekeepers' generative AI offerings in many instances. For example, gatekeepers will remain free to leverage generative AI business user data where the generative AI service is provided separately from other core platform services.¹¹ Similarly, the DMA's prohibition on forcing business and end users to subscribe to or register with any other services of the gatekeeper applies only across core platform services.¹² Without explicit integration of generative AI as a core platform service, gatekeepers remain free to force users to register with their, for example, social networking services as a condition to use their generative AI services.

While the development of large generative AI models is relatively recent in comparison to other core platform services, the potential for some generative AI providers to acquire an entrenched and durable position is significant. At any rate, Article 3(1) allows the Commission to designate a gatekeeper even where an entrenched and durable position is not yet reached but is foreseeable in the near future, which ensures that the Commission does not have to wait for economic power to settle in order to intervene under the DMA. In a similar vein, the quantitative thresholds under Article 3(2) provide a presumption, but the Commission may still designate a gatekeeper where the qualitative conditions are met but the quantitative conditions are not (Articles 3(8) and 17), relieving the Commission of the necessity to wait three years under Article 3(2)(c).

¹⁰The agreement defines AI services as 'services that are labeled or described by Microsoft as including, using, powered by, or being an Artificial Intelligence ("AI") system'.

¹¹See Article 6(2) of the DMA.

¹²See Article 5(8) of the DMA.

For example, in terms of user numbers which are indicative of a service's importance under Article 3(1)(b) and 3(2)(b), ChatGPT is the second fastest digital service to reach 1 million users in just five days (Duarte, 2024) (with an estimated 100 million monthly active users in November 2023), only surpassed by Threads (which inherited Instagram's user base; note that Instagram took 2.5 months to reach 1 million users). Other models, such as Stable Diffusion, acquired a user base of 1 million users in just over a month (Douglas, 2022). Furthermore, other state-of-the-art generative AI models have been developed by companies that already benefit from an entrenched and durable position, making their generative AI services inherit a reputational advantage and most importantly, a large user base—Google's Gemini, Meta's LLaMA, Microsoft's Bing AI (which is powered by GPT-4), among many others. While user numbers, especially business user numbers of these services are not publicly available, the available data demonstrates their rapid and continuous growth.

In terms of market capitalization which is indicative of 'significant impact on the internal market', several generative AI providers have already been found by the Commission to have a significant impact, such as Google, Meta, and Microsoft, who have been developing or integrating large generative AI models in their services. Among other providers of generative AI models, valuations have been rapidly increasing over the last couple of years (Liu, 2023; Vipra and Korinek, 2023). In its latest fundraising round in February 2024, OpenAI was valued at USD 80 billion, which is approximately EUR 73.7 billion—a figure nearly reaching the DMA's quantitative threshold at EUR 75 billion. While there is always a risk of inflation in private valuations, there is little doubt that OpenAI is fast becoming one of the most valuable companies in the world.

Furthermore, there is evidence that unfair practices are emerging in genAI. For example, Microsoft reportedly threatened to cut off two business customers' access to its search engine index data, which it licenses to other search engines, if they did not stop using the data as training input for their own generative AI tools (Nylen and Bass, 2023). Such behaviour appears to be intended to weaken the contestability of Microsoft's own generative AI services by denying potential competitors the key input for model training. Generative AI providers remain free to leverage their superior bargaining power to engage in such unfair practices, and their already considerable economic power to prevent contestability. The DMA provides an opportunity to address contestability and fairness issues in the generative AI industry by designating generative AI as a core platform service. To be clear, including a new core platform service in the DMA does not automatically result in the designation of gatekeepers for that core platform service. A gatekeeper can only be designated if the conditions of Article 3(1) are fulfilled. However, once generative AI becomes a core platform service, the Commission can act swiftly if and once gatekeepers emerge.

Our proposal invites the question of how the relevant generative AI core platform service or services should be defined. There is no need to muddy the regulatory waters: a new core platform service to address generative AI can be defined in reference to the AI Act as 'general-purpose AI systems'. This category was added to the draft AI Act with generative AI applications in mind following the advent of ChatGPT (see Appendix A for more information). This approach would be in keeping with the DMA's definition of core platform services in reference to other EU legislation, to the extent definitions exist in other directives and regulations, e.g. for 'online search engine' and 'cloud computing service'.

An initial assessment of gatekeeper obligations reveals that some of them would already apply to generative AI systems if they were listed in the DMA as a core platform service. For example, under Article 6(2), gatekeepers are prohibited from using, in competition with their business users, the data generated by business users of their core platform services and by these businesses' customers. This provision would prevent generative AI gatekeepers from free-riding on data produced by businesses relying on their API to provide downstream applications, either in generative AI verticals or other industries. Similarly, a generative AI gatekeeper without any other core platform service would be prevented from combining personal data across its services (Article 5(2)). Article 6(13) could be used to ensure that generative AI gatekeepers do not implement disproportionate terms and conditions for terminating the provision of the generative AI service.¹³

5 Conclusion

The DMA should be amended to add generative AI to its list of core platform services. This is the first necessary step to address the emergence of entrenched positions and unfair practices related to gatekeeping in the generative AI industry. We invite the Commission as well as other scholars

¹³We continue to evaluate whether new types of gatekeeper obligations should be added to the DMA specifically addressing generative AI gatekeepers. The paper might be subject to changes accordingly.

and practitioners to further analyse the implications of gatekeeping in generative AI with a view to designing adequate gatekeeper obligations in this context. Furthermore, the facts and data that we presented to make our case were obtained from publicly available resources. They strongly point to the risk of gatekeeping in generative AI, but as we have discussed, opacity as opposed to openness is currently the industry norm. We are confident that the Commission, which has extensive investigative powers, can uncover a great deal more information from generative AI providers demonstrating this risk. While this sector is advancing rapidly, the technological and institutional trends are now clear. It is better to nip gatekeeping in the bud before we find ourselves once again with economic power positions too strong to challenge in any meaningful way.

References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. (2023). GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Adobe (2023). Adobe unveils future of creative cloud with generative AI as a creative co-pilot in Photoshop. Press release. May 23, 2023.
- Amazon (2023). Announcing new tools to help every business embrace generative AI Amazon Web Services. Press release. September 28, 2023.
- Anil, R., Dai, A. M., Firat, O., Johnson, M., Lepikhin, D., Passos, A., Shakeri, S., Taropa, E., Bailey, P., Chen, Z., et al. (2023). Palm 2 technical report. *arXiv preprint arXiv:2305.10403*.
- AnthropicAI (2024). Privacy policy. Press release. March 4, 2024.
- Appenzeller, G., Bornstein, M., and Casado, M. (2023). Navigating the High Cost of AI compute. Web Page. April 27, 2023.
- Arunasalam, S. (2023). ChatGPT users can now browse internet, OpenAI says. Web Page. September 27, 2023.
- Ben-Sasson, H. and Greenberg, R. (2023). 38TB of data accidentally exposed by Microsoft AI researchers. Web Page. September 22, 2023.
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., et al. (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Brockman, G., Sutskever, I., and OpenAI (2019). OpenAI LP. Press release. March 11, 2019.
- Casper, S., Davies, X., Shi, C., Gilbert, T. K., Scheurer, J., Rando, J., Freedman, R., Korbak, T., Lindner, D., Freire, P., et al. (2023). Open problems and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint arXiv:2307.15217*.
- Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H. W., Sutton, C., Gehrmann, S., et al. (2023). Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.
- Coldevey, D. (2023). App Store for AI: Openai’s GPT Store lets you build (and monetize) your own GPT. Web Page. November 7, 2023.
- Council of the EU (2023). Digital Markets Act (DMA): agreement between the Council and the European Parliament. Press release. March 25, 2022.
- Devlin, J., Chang, M., Lee, K., and Toutanova, K. (2019). BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*, pages 4171–4186. Association for Computational Linguistics.
- Devlin, J. and Chang, M.-W. (2018). Open sourcing BERT: State-of-the-art pre-training for natural language processing. Web Page. November 2, 2018.
- Douglas, W. (2022). Generative AI is changing everything. But what’s left when the hype is gone? Technology Review. December 16, 2022.

- Duarte, F. (2024). Number of ChatGPT Users. Web Page. March 1, 2024.
- European Commission (2020a). Impact assessment of the Digital Markets Act. Press release. December 15, 2020.
- European Commission (2020b). Proposal for a regulation of the european parliament and of the council on contestable and fair markets in the digital sector (digital markets act). *SEC (2020) 437 final*.
- European Commission (2020c). Shaping Europe’s Digital Future. https://commission.europa.eu/document/download/84c05739-547a-4b86-9564-76e834dc7a49_en?filename=communication-shaping-europes-digital-future-feb2020_en.pdf. February, 2020.
- European Commission (2022a). Final report—sector inquiry into consumer Internet of Things. COM(2022) 19 final. January 20, 2022.
- European Commission (2022b). On contestable and fair markets in the digital sector and amending Directives (EU) 2019/1937 and (EU) 2020/1828 Digital Markets Act. OJ L 265. December 10, 2022.
- European Commission (2023). Questions and Answers: Digital Markets Act: Ensuring fair and open digital markets. Press release. September 6, 2023.
- European Parliament (2023). AI Act: a step closer to the first rules on artificial intelligence. Press release. May 11, 2023.
- Firm, J. S. L. and Butterick, M. (2023). Github Copilot litigation. January 25, 2024.
- FTC (2023). Generative AI raises competition concerns. June 29, 2023.
- Gemini Team, Anil, R., Borgeaud, S., Wu, Y., Alayrac, J.-B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A. M., Hauth, A., et al. (2023). Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Google (2023). Vertex AI Model Garden and Generative AI Studio. Press release. March 29, 2023.
- Grady, P. and Curnin, C. (2023). Partnering with Harvey: Putting LLMs to Work. Web Page. April 26, 2023.
- Gudibande, A., Wallace, E., Snell, C., Geng, X., Liu, H., Abbeel, P., Levine, S., and Song, D. (2023). The false promise of imitating proprietary LLMs. *arXiv preprint arXiv:2305.15717*.
- Insights, M. T. R. (2023). The great acceleration: CIO perspectives on Generative AI. *MIT & Databricks*.
- Jasper (2023). Jasper vs. ChatGPT: Which is right for you? Web Page.
- Jiang, Z., Lin, H., Zhong, Y., Huang, Q., Chen, Y., Zhang, Z., Peng, Y., Li, X., Xie, C., Nong, S., et al. (2024). MegaScale: Scaling large language model training to more than 10,000 GPUs. *arXiv preprint arXiv:2402.15627*.
- Kanbach, D. K., Heiduk, L., Blueher, G., Schreiter, M., and Lahmann, A. (2023). The GenAI is out of the bottle: generative Artificial Intelligence from a business model innovation perspective. *Review of Managerial Science*, pages 1–32.
- KeyserSosa (2023). An update regarding Reddit’s API. Press release. November 16, 2023.
- Knight, W. (2023). OpenAI’s CEO says the age of giant AI models is already over. Web Page. April 17, 2023.
- Li, H., Su, J., Chen, Y., Li, Q., and ZHANG, Z.-X. (2024). Sheetcopilot: Bringing software productivity to the next level through large language models. *Advances in Neural Information Processing Systems*, 36.
- Lieberman, M. B. and Montgomery, D. B. (1988). First-mover advantages. *Strategic management journal*, 9(S1):41–58.

- Liu, E. (2023). A \$90 billion valuation for OpenAI? Tech's new star is hotter than Nvidia. Web Page. September 28, 2023.
- Lomas, N. (2023). ChatGPT-maker OpenAI accused of string of data protection breaches in GDPR complaint filed by privacy researcher. Web Page. August 30, 2023.
- Macaulay, T. (2023). ChatGPT has generated a brand new app sector—and European devs are leading it. Web Page. May 2, 2023.
- Microsoft (2023a). Introducing deep search. Web Page. December 5, 2023.
- Microsoft (2023b). Microsoft, Anthropic, Google, and OpenAI launch frontier model forum. Web Page. July 26, 2023.
- Microsoft (2023c). Microsoft services agreement. Press release. July 30, 2023.
- Microsoft (2023). Reinventing search with a new AI-powered Microsoft Bing and Edge, your copilot for the web. Press release. May 16, 2023.
- Murgia, M. (2023). Big Tech companies use cloud computing arms to pursue alliances with AI groups. Web Page. February 5, 2023.
- NVIDIA (2021). NVIDIA A100 Tensor Core GPU. Press release. June, 2021.
- Nylen, L. and Bass, D. (2023). Microsoft threatens data restrictions in rival AI search. Web Page. March 25, 2023.
- OpenAI (2019). Better language models and their implications.
- OpenAI (2020). OpenAI API. Press release. June 11, 2020.
- OpenAI (2024a). Enterprise privacy at OpenAI. Press release. January 10, 2024.
- OpenAI (2024b). Europe privacy policy. Press release. December 15, 2023.
- OpenAI (2024c). How ChatGPT and our language models are developed. Press release. January 10, 2024.
- Pinsky, Y. (2023). BARD can now connect to your google apps and services. Web Page. September 19, 2023.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. (2019). Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. (2022). Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3.
- Reid, E. (2023). Supercharging search with generative AI. Press release. May 10, 2023.
- Saparov, A., Pang, R. Y., Padmakumar, V., Joshi, N., Kazemi, M., Kim, N., and He, H. (2024). Testing the general deductive reasoning capacity of large language models using ood examples. *Advances in Neural Information Processing Systems*, 36.
- Schaeffer, R., Miranda, B., and Koyejo, S. (2024). Are emergent abilities of large language models a mirage? *Advances in Neural Information Processing Systems*, 36.
- Sevilla, J., Heim, L., Ho, A., Besiroglu, T., Hobbhahn, M., and Villalobos, P. (2022). Compute trends across three eras of machine learning. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Siemens (2023). Siemens healthineers shows potential of generative AI in medical technology. November 26, 2023.
- Soni, V. (2023). Impact of generative AI on small and medium enterprises' revenue growth: The moderating role of human, technological, and market factors. *Reviews of Contemporary Business Analytics*, 6(1):133–153.

- Thompson, N. C., Greenewald, K., Lee, K., and Manso, G. F. (2021). Deep learning’s diminishing returns: The cost of improvement is becoming unsustainable. *IEEE Spectrum*, 58(10):50–55.
- Thoppilan, R., De Freitas, D., Hall, J., Shazeer, N., Kulshreshtha, A., Cheng, H.-T., Jin, A., Bos, T., Baker, L., Du, Y., et al. (2022). Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Vincent, J. (2019). AI researchers debate the ethics of sharing potentially harmful programs. *Verge*. February 21, 2019.
- Vipra, J. and Korinek, A. (2023). Market concentration implications of foundation models: The invisible hand of ChatGPT. September 7, 2023.
- Vipra, J. and Myers West, S. (2023). Computational power and AI. *AI Now Institute*.
- Weatherbed, J. (2023). OpenAI’s regulatory woes have barely started. *Web Page*. May 5, 2023.
- Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., Yogatama, D., Bosma, M., Zhou, D., Metzler, D., et al. (2022). Emergent abilities of large language models. *Transactions on Machine Learning Research*.
- Wiggers, K. (2023). OpenAI’s Foundry will let customers buy dedicated compute to run its AI models. *Web Page*. February 22, 2023.

A Foundations

A.1 Principles of the Digital Markets Act

Regulation 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector, also known as the ‘Digital Markets Act’, came into force on 1 November 2022. It is the result of a recognition by the EU legislature that existing European competition law rules, established under Articles 101 and 102 of the Treaty on the Functioning of the European Union (TFEU), were insufficient to address forms of economic power associated with certain features of digital markets. In particular, many digital services were thought to have characteristics that had a tendency to result in unfair practices and a lack of contestability. In the run up to its DMA proposal in December 2020, the European Commission (‘Commission’) announced that it would explore ‘**ex ante rules** to ensure that markets characterised by large platforms with significant network effects acting as gatekeepers, remain fair and contestable for innovators, businesses, and new market entrants’, which culminated in the DMA ([European Commission, 2020c](#)).

The DMA is built on two new regulatory categories: ‘core platform services’ and ‘gatekeepers’. ‘Core platform service’ is not defined in the DMA, but a list of core platform services is provided in Article 2(2), comprised of services with a ‘capacity to affect a large number of end users and businesses’ (Recital 14). The DMA outlines numerous characteristics that providers of ‘core platform services’ in the digital sector exploit. These include scale economies, strong network effects, the ability to intermediate between multiple groups of users, high dependence of users on platforms, lock-in effects, and other data-driven advantages (Recitals 2, 13, 25). As a result, some providers of core platform services reach gatekeeping positions. Importantly, such characteristics and their likelihood in resulting in unfair practices and a lack of contestability are thought to exist even where a company or entity does not enjoy a ‘dominant position’—a traditional requirement in EU competition law to apply Article 102 TFEU on abuse of dominance. Article 3 of the DMA, further discussed in the next subsection, sets clear and unambiguous criteria for the designation of gatekeepers. If a company is designated as a gatekeeper, they are automatically obligated to comply with the various rules imposed by the DMA for their relevant core platform services.

Meanwhile, the AI Act is the first comprehensive AI regulation in the world, targeting a broad range of AI applications. The EU legislature reached a political agreement on the AI Act in December 2023, and it is expected to formally become law in the first half of 2024. The AI Act, as proposed

by the Commission, was focused on mitigating risks associated with some applications of AI based on a four-tiered risk framework. Under this framework, which is retained in the latest version of the AI Act, the use of AI with unacceptable risk, such as government use of automated social scoring of its citizens, is banned. The framework establishes specific obligations for AI applications that present high, limited, and minimal risks to health, safety, and fundamental rights of EU citizens. While the high-risk category applications need to conduct conformity assessments to ensure all the obligations are met, the last two categories are left mostly unregulated.

Reflecting on the enormous success of OpenAI's ChatGPT in November 2022, the European Parliament proposed new AI Act provisions that address AI systems which have multiple use cases and hence cannot be easily categorised within the existing four tiers. This category now features in the AI Act after lengthy deliberations with the Council and is called 'general-purpose AI'. While the original proposal from the Commission was focused on determining the risk of different applications based solely on their purpose (for example, applications that were built for the purposes of employment or law enforcement were classified as high risk), following the amendments to the draft, the potential capabilities of general-purpose AI systems will also need to be considered by their providers. The many use cases of generative AI models motivated the MEPs to expand the four-tiered approach. The new section addressing general-purpose AI includes obligations for the providers of foundation models to guarantee 'robust protection of fundamental rights, health and safety and the environment, democracy and the rule of law', as well as a 'need to assess and mitigate risks, comply with design, information and environmental requirements and register in the EU database' ([European Parliament, 2023](#)). Generative AI built on top of foundation models would also have to comply with transparency requirements to mitigate the risk of misinformation from generating 'deep-fakes' in image and video outputs.

The deliberate inclusion of general-purpose AI and foundation models in the AI Act aims to mitigate its distinctive, systemic risks to EU citizens and institutions. It does not, however, address the risks of contestability and any imbalances of bargaining power between generative AI companies and their business users, as this is not within the scope of the AI Act.

A.2 Core Platform Services and Gatekeeper Obligations

The DMA targets 'gatekeeper' companies. A gatekeeper is defined under Article 3(1) as an undertaking that: (i) 'has a significant impact on the [EU's] internal market', (ii) 'provides a core platform service which is an important gateway for business users to reach end-users', and (iii) 'enjoys an entrenched and durable position in its operations, or it is foreseeable that it will enjoy such a position in the near future'. 'Gateway' is broadly construed in the DMA: it covers any platform service used by a business to provide goods or services to any end users. The latter excludes business users.

Article 3(2) provides the corresponding quantitative criteria. If an undertaking meets the quantitative thresholds for any of its core platform services, it is presumed to satisfy the qualitative criteria of Article 3(1) and is therefore designated a 'gatekeeper'. The designation concerns only the core platform services meeting the gatekeeping criteria.

The current list of 'core platform services' provided in Article 2(2) covers ten established digital services.¹⁴ This list of core platform services does not explicitly cover generative AI services. Certain cases of generative AI applications will indirectly fall within the remit of the DMA when gatekeepers integrate proprietary or third-party generative AI into their core platform services, such as search engines. However, the DMA's gatekeeper obligations (Articles 5, 6 and 7) designed to counteract contestability and fairness issues will not apply to their generative AI offering when it is provided as a standalone platform service. In the remainder of this paper, we present ways in which generative AI is provided as a platform service and argue for its explicit integration into the DMA.

Nearly two years passed between the Commission's DMA proposal in December 2020 and its entry into force (this was relatively quick for the EU's law-making process). The European Parliament and the Council of the European Union, the EU's other two legislative bodies, decided their positions largely in 2021. The DMA reached its near-final shape in March 2022 ([Council of the EU, 2023](#)).

¹⁴Such as online intermediation services; online search engines; online social networking services; video-sharing platform services; number-independent interpersonal communications services; operating systems; web browsers; virtual assistants; cloud computing services; and online advertising services. The first round of designations by the Commission includes six gatekeepers and eight different categories of core platform services.

Generative AI models have been gaining traction since 2018, but arguably it was not until the launch of ChatGPT in late 2022 that the public at large and policymakers became aware of the significance of generative AI. By then, the DMA was already in place.

The drafters of the regulation foresaw that the DMA would need to be amended with technological developments and put in place mechanisms to keep the regulation up to date. Recital 77 states: ‘The services in the digital sector and the types of practices relating to these services can change quickly and to a significant extent’, which has been the case for generative AI services. Article 19 provides the procedure that the Commission needs to follow to add/drop core platform services. We explain this process below and lay out the conditions under which a new core platform service should be added to the DMA.

A.3 Expanding the DMA’s List of Core Platform Services

The DMA does not define ‘core platform service’. Article 2 (‘Definitions’) merely provides the list of core platform services without a specific definition. Likewise, Article 19 gives a lot of leeway to the Commission in its examination of potential new core platform services and does not provide any specific criteria. The Commission, in its press release concerning the first designation of six gatekeepers, offered the following definition:

Core platform services are those services in the digital economy that exhibit certain features and where absent regulatory intervention the identified failures would effectively remain un-addressed. Such features entail highly concentrated services, where usually one or very few large digital platforms set the commercial conditions with considerable autonomy and where few large digital platforms act as gateways for business users to reach their customers. ([European Commission, 2023](#))

Furthermore, while DMA does not provide a functional definition of core platform services, it does delineate characteristics of the core platform services that it aims to address (Recitals 2 and 14). These characteristics include ‘extreme scale economies’, ‘very strong network effects’, ‘a significant degree of dependence of both business users and end users’, ‘lock-in effects’, ‘a lack of multi-homing’, ‘vertical integration’, and ‘data driven-advantages’ (Recitals 2 and 14). As such, the DMA arguably conflates ‘core platform service’ and ‘gatekeeper’: many of these characteristics are also listed under Article 3(8) concerning Commission decisions to designate a gatekeeper where the quantitative thresholds are not met.

The DMA clearly outlines a mechanism to add new or remove existing core platform services in Article 19. For the purpose of examining whether one (or more) digital services should be added to the list of core platform services, the Commission can open a market investigation. The Commission must then publish a report within 18 months and submit a legislative proposal to the European Parliament and the Council. From there, the amendment enters the ordinary legislative procedure. Adding new core platform services to the DMA is ‘a core part of keeping the regulation future-proof’ ([European Commission, 2020a](#)). The dynamic nature of digital markets requires that the Commission maintains regular, continuous monitoring of its impact. Therefore, there are no legislative barriers to adding new core platform services to the DMA.

The expansion of the list of core platform services has already transpired during the inter-institutional negotiations. At first, the DMA proposal ([European Commission, 2020b](#)) listed eight out of the current ten core platform services: online intermediation services, online search engines, social networking, video sharing platform services, number-independent interpersonal electronic communication services, operation systems, cloud services, and advertising services. Subsequently, the final text ([European Commission, 2022b](#)) included two additional core platform services: web browsers and virtual assistants. The decision to include virtual assistants was based on the Commission’s sector inquiry into the consumer Internet of Things (IoT) ([European Commission, 2022a](#)), which found competition concerns such as high entry barriers and difficulties in competing with companies that are vertically integrated and form their own ecosystems.

Based on the DMA, the Impact Assessment, and the Commission’s press releases, we propose that the following three cumulative criteria need to be met for the Commission to designate a new core platform service.

First, the prospective digital service should constitute a **gateway for business users to reach end users**. The DMA excludes platforms that have no role in bringing together business users with end users. While multi-sided markets such as search engines, video-sharing platforms, and operating systems are canonical examples of gateways that directly intermediate transactions between businesses and consumers, the concept of gateway is construed more broadly than match-making in the DMA. For example, it designates cloud computing service as a core platform service because although it does not match-make transactions, it provides the infrastructure to support and enable business users to reach customers.

Second, the service or the market in which it operates must exhibit certain problematic **characteristics that facilitate the emergence of gatekeepers**. Drawing on the DMA and the Impact Assessment, these characteristics include strong network effects which make markets prone to ‘tipping’, extreme scale economies, vertical integration, data-driven advantages, and lack of multi-homing. These characteristics are largely drawn from past antitrust cases, sector inquiries, and academic (especially economic) literature.

Finally, at the risk of stating the obvious, the new core platform service should be **sufficiently distinguishable from existing core platform services**. This, however, does not exclude the possibility of distinguishing between different core platform services even when they are, at times, provided in an integrated fashion. For example, Google Shopping is provided both as a distinct service and within Google Search. Alphabet was designated as gatekeeper separately for these two core platform services.