

cerre

Centre on Regulation in Europe



IS THE DMA READY FOR AGENTIC AI?

REPORT

July 2025

Friso Bostoen
Jan Krämer



The Cerre logo consists of a solid blue square. Inside the square, the word "cerre" is written in a white, lowercase, sans-serif font, centered horizontally and vertically.

cerre

Report

Is the DMA Ready for Agentic AI?

Friso Bostoen
Jan Krämer

July 2025



As provided for in CERRE's bylaws and procedural rules from its “Transparency & Independence Policy”, all CERRE research projects and reports are completed in accordance with the strictest academic independence.

The project, within the framework of which this report has been prepared, received the support and/or input of the following CERRE member organisations: Arcep, EETT, Meta, Ofcom, and Qualcomm. However, they bear no responsibility for the contents of this report. The views expressed in this CERRE report are attributable only to the authors in a personal capacity and not to any institution with which they are associated. In addition, they do not necessarily correspond either to those of CERRE, or of any sponsor or of members of CERRE.

© Copyright 2025, Centre on Regulation in Europe (CERRE)

info@cerre.eu – www.cerre.eu



Table of Contents

<u>ABOUT CERRE.....</u>	<u>4</u>
<u>ABOUT THE AUTHORS</u>	<u>5</u>
<u>EXECUTIVE SUMMARY</u>	<u>6</u>
<u>1. INTRODUCTION</u>	<u>9</u>
<u>2. BUILDING BLOCKS OF AGENTIC AI</u>	<u>12</u>
<u>3. POTENTIAL COMPETITION CONCERNS.....</u>	<u>17</u>
3.1. FORECLOSURE OF AN AI AGENT	17
INPUT FORECLOSURE	17
DISTRIBUTION FORECLOSURE	19
3.2. FORECLOSURE BY AN AI AGENT	20
<u>4. IS AGENTIC AI WITHIN THE DMA'S SCOPE?.....</u>	<u>22</u>
4.1. DESIGNATION AS CORE PLATFORM SERVICE	22
ONLINE INTERMEDIATION SERVICE	23
ONLINE SEARCH ENGINE	23
OPERATING SYSTEM.....	25
WEB BROWSER	25
VIRTUAL ASSISTANT	26
CLOUD COMPUTING SERVICE.....	27
4.2. DETERMINATION OF GATEKEEPER STATUS	28
4.3. ANALYSIS	29
<u>5. ARE THE DMA'S OBLIGATIONS MEANINGFUL IN THE CONTEXT OF AGENTIC AI?</u>	<u>32</u>
5.1. OBLIGATIONS FOR OSS	32
UNINSTALLATION.....	32
INSTALLATION	33
5.2. OBLIGATIONS FOR OSS, VIRTUAL ASSISTANTS, WEB BROWSERS.....	33
DEFAULT SETTINGS AND CHOICE SCREENS	34
INTEROPERABILITY	35
5.3. OBLIGATIONS FOR SEARCH ENGINES	36
SELF-PREFERENCING IN RANKING	37
DATA SHARING	37



FRAND TERMS..... 38

5.4. OBLIGATIONS INDEPENDENT OF CPS CATEGORY 39

RESTRICTIONS ON DATA COMBINATION AND USE 39

DATA PORTABILITY 40

CONTRACTUAL OBLIGATIONS AND RESTRICTIONS..... 41

6. SHOULD THE DMA BE REVISED IN THE FACE OF AGENTIC AI?..... 43

6.1. FORECLOSURE OF AN AI AGENT 43

6.2. FORECLOSURE BY AN AI AGENT..... 46

7. CONCLUSIONS AND RECOMMENDATIONS 50

REFERENCES 53



About CERRE

Providing high quality studies and dissemination activities, the Centre on Regulation in Europe (CERRE) is a not-for-profit think tank. It promotes robust and consistent regulation in Europe's network and digital industry and service sectors as well as in those impacted by the digital and energy transitions. CERRE's members are regulatory authorities and companies operating in these sectors, as well as universities.

CERRE's added value is based on:

- its original, multidisciplinary and cross-sector approach covering a variety of markets, e.g., energy, mobility, sustainability, tech, media, telecom, etc.;
- the widely acknowledged academic credentials and policy experience of its research team and associated staff members;
- its scientific independence and impartiality; and,
- the direct relevance and timeliness of its contributions to the policy and regulatory development process impacting network industry players and the markets for their goods and services.

CERRE's activities include contributions to the development of norms, standards, and policy recommendations related to the regulation of service providers, to the specification of market rules and to improvements in the management of infrastructure in a changing political, economic, technological, and social environment. CERRE's work also aims to clarify the respective roles of market operators, governments, and regulatory authorities, as well as contribute to the enhancement of those organisations' expertise in addressing regulatory issues of relevance to their activities.



About the Authors



Friso Bostoen is Assistant Professor of Competition Law and Digital Regulation at Tilburg University. Previously, he was a Max Weber Fellow at the European University Institute. He holds degrees from KU Leuven (PhD, LLM) and Harvard University (LLM). Friso's research focuses on competition policy for digital markets. His work has resulted in numerous international publications, presentations, and awards (including the AdC Competition Policy Award 2019 and the Concurrences PhD Award 2022).



Jan Krämer is a CERRE Academic Co-Director and Professor for Information Systems at the University of Passau, Germany, where he holds the Chair for Internet & Telecommunications Business. He has a diploma degree in Business Engineering and Management, and a Ph.D. in Economics, both from the Karlsruhe Institute of Technology. He has published numerous articles in the leading journals in the areas of Information Systems, Economics and Marketing. His current research interests include the regulation of platforms and internet markets, as well as digital ecosystems and data-driven business models.



Executive Summary

The future of AI is agentic. We can discern the first signs of agentic AI in the widely used assistants (e.g., ChatGPT) and, more clearly, in the prototypes of AI agents (e.g. Operator). AI agents do not simply respond to user inputs but exhibit proactive, autonomous behaviour based on a user-set prompt or goal – or even based on a user need inferred by the AI agent. In fulfilling tasks, AI agents adapt and learn from their experiences, interact with the wider digital and physical environment, rely on persistent memory, and integrate with external tools and applications (including other AI agents).

Agentic AI has the potential to transform consumer behaviour and, accordingly, to disrupt existing gatekeepers in platform markets from operating systems (OSs) to search and e-commerce. There are, however, two competition risks: one more immediate, the other more distant. The more immediate risk is that incumbents in markets that AI agents rely on limit the growth of third-party AI agents in favour of their own agents and/or to safeguard their position in markets threatened by agentic AI (“foreclosure of AI agents”). The more distant risk is that AI agents become gatekeepers in their own right and use that position to anticompetitively steer demand (“foreclosure by AI agents”).

This report systematically discusses the appropriate regulatory response to these two risks. While reference is made to competition law, our focus is on the Digital Markets Act (DMA), which is likely the best contender for addressing contestability issues related to consumer-facing digital services in the context of AI agents. However, the DMA was not adopted with AI in mind, and while competition authorities have published reports on generative AI/foundation models, they have not studied AI agents. In other words, despite the (expected) rise of AI agents, there is currently little regulatory guidance, which is a gap that this report fills.

The questions tackled in the report are the following: (i) what are the potential competition concerns (foreclosure *of* vs *by* AI agents); (ii) do the firms at the source of these concerns fall under the scope of the DMA; (iii) are the DMA’s obligation meaningful in this context? Based on the answers to these questions, we make recommendations to ensure that the DMA remains an adequate regulatory response in case the potential competition risks manifest themselves. In answering the above questions, we take into consideration several policy interests, including regulatory effectiveness, legal (and investment) certainty, and political legitimacy – and do so amidst a transatlantic debate on the desirability of intervention in (nascent) digital markets.

The **foreclosure of AI agents** could have several sources. First, AI agents rely on foundation models as an input. Agent providers can currently rely on third-party models that are made available in open-source or via API calls. Should such access be degraded in the future, however, building one’s own state-of-the-art foundation model is – especially given data requirements – a difficult proposition. Second, AI agents need to “live” within a (mobile) device, and need access to its hardware and software features to function effectively. Device makers with their own agent may want to advantage their own agent (e.g. via pre-installation and default status) and/or disadvantage those of competitors (e.g. via decreased interoperability).

The scope of the DMA presents no significant gaps when it comes to the foreclosure of AI agents. The most likely chokepoint is the (mobile) OS. OSs qualify as core platform service (CPS) and the most relevant ones (iOS, Android, Windows) have been designated with gatekeeper status. AI foundation



models do not qualify as CPS while one of their key inputs, cloud services (for compute), do. Given the current state of the market, this coverage appears sufficient.

The DMA imposes a number of obligations on gatekeepers that can help prevent the foreclosure of AI agents by keeping their OSs open. First, OS gatekeepers must allow users to uninstall a pre-installed agent as long as it is not essential (Article 6(3), para 1). Second, if the OS gatekeeper's own AI agent is designated as virtual assistant, web browser or search engine, it must go a step further and show an AI agent choice screen (Article 6(3), para 2). Third, the OS gatekeeper must provide third-party AI agents the same interoperability with hardware and software features as is available to its own agent (Article 6(7)). Fourth, if the agent itself is designated as CPS with gatekeeper status (an issue to which we return below), data portability must be guaranteed (Article 6(9)–(10)).

To different degrees, these obligations – along with some others – can help prevent input foreclosure (through denial of equal access to data and on-device hardware and software features) and distribution foreclosure (through pre-installation, self-preferencing and tying). A major challenge will be technical: to function effectively, AI agents require deep integration with the OS, which makes interoperability and the replacement of first-party AI agents complex.

To ensure the DMA adequately covers foreclosure of AI agents, we make two recommendations. The first is to extend the click and query data-sharing obligations of Article 6(11) to include virtual assistants. Given the critical role of such data for AI agent contestability, this measure would enhance market openness. The second recommendation is to add virtual assistants alongside browser engines and other services in Article 5(7), which prohibits gatekeepers from forcing business users to integrate with certain ancillary services. This would help ensure that users retain freedom in their choice of AI agents and, as a result, foster competitive dynamics between such agents. A possible side-recommendation is to extend the CPSs covered by the FRAND access obligation of Article 6(12) to include cloud computing services, which could help ensure non-discriminatory access as the market sees increasing vertical integration.

In the future, **foreclosure by an AI agent** could arise. This depends on the development trajectory of AI agents. It may tend towards concentration due to market features (e.g. the importance of learning from users) or conduct of players in control of key inputs and distribution channels (e.g. those in control of device hardware and software). Should only a few agents remain, or more remain but users are nevertheless “locked in” to their AI agent, then agent providers could use their demand-steering power to foreclose. Such foreclosure generally takes the form of leveraging, in which the AI agent provider props up their other services (or affiliated services) via tying or self-preferencing.

The scope of the DMA becomes important when it comes to foreclosure by an AI agent. In this scenario, the agent itself is the chokepoint, so the question is whether it is subject to the DMA. Given the DMA's wide CPS definitions, the first generations of AI agents may qualify as search engines or web browsers; in the future, a qualification as OS may even become appropriate. Designation in the virtual assistant CPS category is, however, most suitable. The definition (the processing of “demands, tasks or questions”) can accommodate AI agents, even if it was adopted with voice assistants in mind.

Nevertheless, we recommend mitigating potential ambiguities in (and around) the definition of “virtual assistant”. First, the DMA's Annex, which identifies business users of virtual assistants as developers that make their app accessible via the assistant, is restrictive: some developers may



actively make their apps accessible via AI agents, but others are passively “called upon” by the agent. Hence, an amendment of the Annex via delegated act is recommended. Second, as the distance between AI agents and the virtual assistant definition grows, e.g. due to their autonomy (reacting to user “needs” rather than input), legal certainty would benefit from legislatively adapting/replacing the definition. This has the added benefit of giving the inclusion of AI in the DMA political legitimacy.

A number of DMA obligations can help prevent foreclosure by AI agents, though their applicability depends on the CPS category in which the agent is designated. If the agent is designated in either the OS, virtual assistant or web browser CPS category, and it defaults to another service, the gatekeeper must make that default setting changeable; in addition, it must show a search engine choice screen (Article 6(3), para 2). If the agent is designated as search engine or virtual assistant, the gatekeeper cannot self-preference in ranking (Article 6(5)). As an OS or virtual assistant, the AI agent must grant equal interoperability, e.g. with connected devices (Article 6(7)). A number of other obligations related to data and contractual restrictions apply independent of CPS category.

In our view, the most fitting qualification as virtual assistant leads to correct coverage in terms of obligations to prevent foreclosure by an AI agent. Hence, assuming AI agents are designated within the virtual assistant category (or in a future AI agent category), concerns about foreclosure by AI agents do not currently warrant amendments of the DMA.

In conclusion, the DMA is surprisingly ready for agentic AI. The DMA’s unexpected future-proofness in this context stems from two facts: (i) the source of foreclosure of AI agents is likely to be found in CPSs that are already designated, particularly (mobile) OSs; and (ii) legislators added a “virtual assistant” category to the DMA’s list of CPSs, with a wide definition that can plausibly accommodate AI agents. Some amendments, in particular to the definitions of business users and – over time – virtual assistants, as well as the obligations in Articles 5(7), 6(11) and 6(12), could help safeguard the DMA’s future-proofness.



1. Introduction

Agentic AI (or AI agents) is widely regarded as the next evolution of AI, poised to transform how we interact with and utilise technology. These autonomous AI systems possess the capability of independent action and decision-making. Unlike contemporary AI assistants (such as ChatGPT, Copilot, or chatbots), which primarily respond to user inputs and operate within chat-based environments, agentic AI systems exhibit proactive behaviour. Following an initial prompt, they can continue their work without further input or human intervention, and may even plan ahead and execute actions autonomously. In contrast to traditional AI, agentic AI systems possess persistent memory, enabling them to adapt and learn from their experiences. Additionally, they possess enhanced connectivity, allowing them to accept (multi-modal) input from various sources (e.g. voice, text, vision, sensors, or external applications) to effectively comprehend and carry out tasks. Furthermore, agentic AI systems seamlessly integrate with external applications and tools, facilitating even communication, invocation, and collaboration with other AI agents.

While this may sound like science fiction to some, the first prototypes of agentic AI systems have already been showcased by players like OpenAI (called “Operator”),¹ Google (DeepMind’s “Project Mariner”),² Anthropic (Claude’s “Computer Use”),³ and Perplexity (its general “Assistant” and its specialised shopping agent “Buy with Pro”).⁴ In parallel, AI agent startups are experiencing a significant increase in investor funding.⁵ Currently, the prototypes of AI agents are based on a “act-like-a-human” architecture, where the AI agent can “see” what the user typically sees (by capturing the user’s interface, e.g. the browser window or the entire user’s screen) and can perform actions through clicking or typing just like a user would. By mimicking human interaction, the AI can make use of existing interfaces designed for humans. In turn, this does not require application programming interfaces (APIs) specifically designed for AI agents, or additional permissions and entitlements. However, this is likely only a first step in the evolution of AI agents. In the future, AI agents will become more deeply imbedded into technology, and technology and interfaces are likely to become more agent-centric. As AI agents advance, and their use becomes more widespread, new paradigms and forms of interaction will emerge, which have the potential to fundamentally change the technological landscape. This revolution has the potential to redefine industries, disrupt traditional business models, and create a new era of human–machine collaboration.

In particular, AI may radically transform how consumers interact with their (mobile) devices, potentially leading to new interface designs, and new types of wearable devices, where the AI agent, rather than a specific device, app or search engine, provides the unique access point for users to discover and use digital services. This has the potential to disrupt the “gatekeepers” that are currently controlling today’s access points like search, app stores or operating systems. First signs of such new designs of mobile devices are already visible today.⁶ For example, in the vision of an “app-less” AI

¹ <https://openai.com/index/introducing-operator/>.

² <https://deepmind.google/technologies/project-mariner/>.

³ <https://www.anthropic.com/news/3-5-models-and-computer-use>.

⁴ <https://www.perplexity.ai/help-center/en/articles/10450852-basic-information-about-the-perplexity-assistant>,
<https://www.perplexity.ai/hub/blog/shop-like-a-pro>.

⁵ See Stefan Hunt, Emily Chissell and Aman Mawar, ‘Will 2025 be the year of the agent? A primer for competition practitioners on the next wave of AI innovation’ (2025) 9 Competition Law & Policy Debate 20.

⁶ OpenAI is also working on a new wearable device, see <https://openai.com/sam-and-jony/>.



phone,⁷ the user gives tasks (e.g. to book a business trip) to the AI agent, and it is the AI agent – rather than the user – that chooses which underlying apps or providers are employed to fulfil that task. The power to steer consumers' demand, and even to act on behalf of the user, thus lies with the AI agent – potentially creating a new “super-gatekeeper”.

Competition authorities and policymakers have already been active in monitoring and scrutinising potential competition and contestability issues related to AI, specifically with a focus on foundation models, as attested by reports by the CMA⁸ and Autorité de la Concurrence.⁹ CERRE has also published an issue paper focusing on cloud and AI models.¹⁰ But much less attention has been devoted so far to the issue of agentic AI,¹¹ which is situated further downstream and has unique competitive dynamics. In our preceding issue paper, we highlighted potential contestability and competition concerns related to AI agents in mobile ecosystems, and provided an overview on whether and how the EU's digital laws may be able to address some of these issues.¹² We concluded that the Digital Markets Act (DMA) is likely the best contender for addressing any competition issues related to consumer-facing digital services in the context of AI agents. At the same time, the DMA was clearly not drafted with the context of AI (agents) in mind, which can create issues of legal certainty and accordingly pose risks for the investment environment in the EU. Moreover, the Draghi report has rekindled concerns about the costs of precautionary (and in particular symmetric) regulation, even though it also calls for effective enforcement of the (asymmetric) DMA.¹³

Against this background, this report zooms in on the relationship between AI agents and the DMA, and takes our previous analysis one step further by exploring in much more detail whether the DMA, the EU's landmark regulation for addressing competition issues in digital markets, is fit for the advent of agentic AI. Specifically, we analyse whether agentic AI could be subsumed under any of the existing core platform service (CPS) categories, and whether the ensuing obligations are meaningful in this context.

We conclude that the legal text of the DMA can be interpreted in such a way that it is surprisingly well-prepared for the new phenomenon of AI agents. The CPS category of “virtual assistants”, in particular, seems to allow for the potential designation of AI agents under the DMA, at least before they reach a high level of autonomy. Although lawmakers' intent was to capture mainly “voice assistants” (like Alexa or Google Home), the legal text also allows for the designation of some AI agents. In this way, the DMA can already, in principle, address many of the potential competition concerns related to both the foreclosure of an emerging AI agent (i.e. the foreclosure by an existing OS gatekeeper with pre-installed own AI agent of another AI agent), as well as more future-looking, the foreclosure of other services by a DMA-designated AI agent. However, we also make recommendations on select amendments to the DMA and useful clarifications of the interpretation of the legal text that make the DMA even more future-proof in the context of agentic AI.

⁷ See <https://www.telekom.com/en/media/media-information/archive/deutsche-telekom-frees-smartphones-from-apps-1060272>.

⁸ Competition & Markets Authority, ‘AI Foundation Models’ (Initial Report, September 2023) and Competition & Markets Authority, ‘AI Foundation Models’ (Update Paper, April 2024).

⁹ Autorité de la Concurrence, Opinion 24-A-05 of 28 June 2024 on the competitive functioning of the generative artificial intelligence sector.

¹⁰ Zach Meyers and Marc Bourreau, ‘A Competition Policy for Cloud and AI’ (CERRE Issue Paper, 2025).

¹¹ For an exception, see Hunt, Chissell Mawar (n 5).

¹² Friso Bostoen and Jan Krämer, ‘AI Agents and Ecosystems Contestability’ (CERRE Issue Paper, 2024).

¹³ Mario Draghi, ‘The future of European competitiveness Part A | A competitiveness strategy for Europe’ (Report, 2024) 30; Mario Draghi, ‘The future of European competitiveness Part B | In-depth analysis and recommendations’ (Report, 2024) 244 and 302–304.



The remainder of this paper is organised as follows. In Section 2 we provide technical background on the definition and main building blocks of AI agents (or agentic AI). In Section 3 we summarise the main potential competition concerns that arise in the context of agentic AI. They fall broadly in two categories: foreclosure *of* an (emerging) AI agent and foreclosure *by* a (gatekeeping) AI agent. In Section 4 we analyse whether and in which CPS category AI Agents could potentially be designated under the DMA. In Section 5 we interpret the legal text of the DMA in the context of AI agents and in light of various designation possibilities. In Section 6, we then discuss whether the DMA is fit to address the competition concerns related to a foreclose of AI agents and foreclosure by an AI agent, respectively. In Section 7, finally, we conclude by making policy recommendations in aspects in which the DMA can be improved with regards to its interpretation and fit in the context of agentic AI.



2. Building Blocks of Agentic AI

Agentic AI has recently become a buzzword and is not yet well defined.¹⁴ In the following, we offer our understanding of agentic AI, and its main building blocks, which we rely on for the remainder of the report.

Agentic AI is distinct from “generative AI” (GenAI), although it likely employs GenAI. The defining characteristic of agentic AI is its **greater degree of autonomy** compared to other AI forms. It can and will make decisions independently with minimal human oversight. The extent of autonomy and human oversight involved will likely evolve gradually and be task-specific. Currently, AI generally operates based on human-defined objectives. Fully autonomous AI systems that can set and pursue their own goals remain theoretical. While many AI systems exhibit some agentic behaviours, they are typically specialised for specific tasks and limited in scope due to safety and usability concerns. Agentic general AI capable of autonomous decision-making across a broader and more diverse range of tasks does not exist yet. The pursuit of such capabilities has sparked both interest and concern among researchers. Some experts raise significant safety concerns regarding potential risks, while others view agentic AI as the ultimate goal towards artificial general intelligence (AGI). In essence, what distinguishes agentic AI from GenAI and other forms of AI (AI itself being a poorly defined and often loosely used term) is constantly evolving, contributing to the perception of the term “agentic AI” as fuzzy.

Figure 1: Distinctions between agentic AI, GenAI and traditional AI according to AISERA¹⁵

Feature	Agentic AI	Generative AI	Traditional AI
Primary Function	Goal-oriented action & decision-making	Content generation (text, code, images, etc.)	Focused on automating repetitive tasks
Autonomy	High – Operates with minimal human oversight	Variable – May require user prompts or guidance	Low – Relies on specific algorithms and set rules
Learning	Reinforced Learning – Improves through experience	Data-driven learning – Learns from existing data	Relies on predefined rules and human intervention

For the purpose of this report, it is useful to differentiate AI agents from AI assistants, like ChatGPT. AI assistants based on large language models (LLMs) are by now very familiar to us. AI agents go much beyond the functionalities currently associated with AI assistants.¹⁶ They demonstrate not only greater autonomy but also the ability to perceive their surroundings, plan actions, and attain specific objectives. Moreover, they **adapt and learn from their experiences**. While GenAI is mainly built on pre-trained models (that learned from existing data), agentic AI is associated more with learning from experience and persistent memory. Some scholars have argued that data-driven network effects (positive data feedback loops) – such as associated with search engines, where users provide implicit feedback by clicking and viewing search results – are not strong in the context of GenAI or AI assistants,

¹⁴ See, e.g., <https://techcrunch.com/2025/03/14/no-one-knows-what-the-hell-an-ai-agent-is/>.

¹⁵ <https://aisera.com/blog/agentic-ai/>.

¹⁶ We largely follow the definition of IBM here, see <https://www.ibm.com/think/topics/ai-agents-vs-ai-assistants>.



precisely because this technology is using pre-trained models and does not learn from user input.¹⁷ Hence, a tipping of these markets, like in the case of the search engine market is less likely in the context of markets for AI. However, this reasoning does not seem as relevant in the context of agentic AI, and indeed newer advancements in GenAI, which have more persistent memory¹⁸, and rely more heavily on reinforcement learning, a machine learning technique where the model is improved through a carrots-and-sticks approach. Positive outcomes are positively reinforced (rewarded) and negative outcomes are negatively reinforced (punished). Indeed, newer advancements in the quality of GenAI assistants, which can now perform complex “reasoning” (e.g. ChatGPT models ending in “o”) have been achieved through reinforcement learning.¹⁹ In order to perform reinforcement learning, user feedback is required on what constitutes “good” outcomes that shall be reinforced. In order to scale this approach, typically one AI model is trained (based on user feedback) to be the judge of the outcomes provided by another AI model, which is to be improved.²⁰ Evidently, the more user feedback is available, the better is the judging model aligned with the users’ underlying preferences,²¹ and hence the better is the quality of the outcomes provided by the AI.

Further, in order to turn an LLM into an assistant that responds to user requests, large amounts of conversational datasets, which serve as a model for conversations between humans and agents, are required. Conversational datasets can thus be very important for the quality of AI assistants and agents. While there are open-source conversational datasets available (e.g. derived from movies), creating and curating new conversational data sets remain a pivotal source of advancing the quality of LLMs and are hence a source of competitive advantage for GenAI.²² For example, available conversational datasets may not be representative of the specific domain or use case that the model is being trained for. Thus, being able to derive conversational datasets from the interactions between users and the agent can prove a valuable source for improving the agent’s quality.

Taken together, it is therefore reasonable to think that data-driven network effects are very relevant in the context of Agentic AI. This is even more so true for agentic AI systems that serve as personal assistants, say on a smart device, fulfilling tasks that traditionally a human personal assistant would do, such as booking of travel itineraries. In this context, the AI system, like a human assistant, needs to know the idiosyncratic tastes and preferences of its user. User feedback data is thus a vital input to agentic AI, more so than for non-agentic AI.

Next to learning, another vital building block for agentic AI systems is its ability to **perceive and interact with its environment**. While AI assistants, as we know them, predominantly live in a browser or app window, and interact with the user predominantly through text, voice or image input and output, agentic AI systems are characterised by their ability to sense and interpret environmental data (e.g. through sensors), and to affect and act in the environment through so-called “tools”. Tools are typically software components that enable the AI to perform certain actions. This is not limited to the virtual environment (e.g. moving the mouse, taking a screenshot), but can also result in actions in the real world (e.g. turning the lights on, commanding a robot, booking a flight). In order to perceive and

¹⁷ Compare Andrei Hagiu and Julian Wright, ‘Artificial intelligence and competition policy’ (2025) International Journal of Industrial Organization (advance access).

¹⁸ See <https://openai.com/index/memory-and-new-controls-for-chatgpt/>.

¹⁹ See, e.g., <https://openai.com/index/learning-to-reason-with-llms/>.

²⁰ See Daniel Ziegler et al, ‘Fine-tuning language models from human preferences’ (2020) arXiv:1909.08593.

²¹ Although there are likely diminishing returns to more data when the dataset is large enough, see Jan Krämer and Daniel Schnurr, ‘Big data and digital markets contestability: Theory of harm and data access remedies’ (2022) 18 Journal of Competition Law & Economics 255.

²² Compare <https://defined.ai/blog/open-source-datasets-for-conversational-ai-advantages-and-limitations>.



interact with its environment, agentic AI needs access to a larger array of inputs and to tools. In practice, this requires **deep integration of the agentic AI system with the operating system (OS) of the device in which the AI lives**. For example, AI agents may require access to the device's GPS location, calendar, maps and browser to perform the tasks of an automated travel booking. AI agents may further need access to low-level OS functionalities, such as power management in order to run efficiently in the background, or to be able react to handsfree voice commands. Further, AI agents require access to tools provided by external content and service providers, such as mobility providers, search engines or payment providers in order to fulfil their tasks, such as searching, booking and paying a travel itinerary. These tools need not necessarily be downloaded and reside as software applications on the phone. Instead, the information exchange between the AI agent and the application could be handled purely through appropriate APIs.

The prototypes of agentic AI systems released by OpenAI, Google and Anthropic all lack such deep integration with the OS and therefore rely on mimicking user input in the browser. Browsers in itself are already benefitting from deeper OS integration than other apps, and by their nature, can access all web content. Therefore, browsers already open up a window for interacting with the environment that does not require further integration or permissions by the OS. In the future, however, integration must and will go deeper. As we highlight in more detail below, in the future agentic AI and the OS may even integrate to the extent that they cannot be separated anymore.

Another key component of agentic AI systems is their ability to **plan and reason** (based on their perceived inputs). Planning and reasoning capabilities are built on the foundation of AI models. These come in two flavours. First, general-purpose **foundation models** such as those pre-trained transformer models popularised by OpenAI (ChatGPT), Meta (Llama), Google (Gemini) or Anthropic (Claude). Foundation models (which include LLMs) are large-scale, general-purpose AI models trained on broad and diverse datasets – practically all information accessible on the web. They are designed to perform well across many different tasks without needing major changes to their architecture. Training such foundation models is very cost-intensive and the market for the provision of such foundation models may end up being highly concentrated.²³ Second, **specialised models** which are finetuned or built specifically for a narrow task or domain. They may be built from scratch but are often based on foundation models optimised for performance in a specific area. An AI agent will typically only make use of one foundation model, but resort to several specialised models for specific tasks, such as shopping online, or scheduling events. In fact, some note that in the future even an **ecosystem of AI agents** may emerge, where more general-purpose agents break up certain tasks and complex workflows into smaller subtasks, and then cooperate with or rely on more specialised agents to fulfil some of the (sub-)tasks.²⁴

While both foundation models and specialised models can be provided for the AI agent, fully or in part by third parties, the core component of an AI agent is the **agentic framework**, which puts the various building blocks together. This is what makes AI “agentic”: going beyond passive chat to take actions and plan sequences toward goals autonomously. This includes task planning, memory and state management, goal prioritisation and execution, and which models to use for what task. It also includes

²³ Competition & Markets Authority, Initial Report and Update Paper (n 8) and Autorité de la concurrence (n 9). See also Klaus Kowalski, Cristina Volpin and Zsolt Zombori, ‘Competition in Generative AI and Virtual Worlds’, European Commission Competition Policy Brief 3/2024. On the concentrating and deconcentrating forces in this market, and where they may lead, see Meyers and Bourreau (n 10).

²⁴ See, e.g., <https://www.assistants.ai/blog/agentic-ai-components>.



the integration code for tool access, i.e. a program that translates the AI agent's abstract tool requests (like "move mouse" or "take screenshot") into actual operations in the virtual environment, and the "agent loop", i.e. a program that handles communication between the agent and the environment, sending actions to the environment and returning the results (screenshots, command outputs) back to the agent.²⁵ However, given the nascent state of agentic AI, there are few industry standards on what comprises an agentic framework and how agents can interact with the environment.²⁶ This likely to undergo major developments in the future. Ultimately, this is the part that defines the agent and on the basis of which agents compete that cannot be sourced from a third-party.

Finally, AI agents require access to **computing resources**, both for training and running. For AI agents running on mobile devices, *mobile AI chips* capable of running AI models at the edge (i.e. on device) are important to increase user experience, as they reduce latency and dependency on connectivity compared to cloud-based solutions. Furthermore, AI computing "on chip" can improve privacy as sensitive data does not have to leave the device.²⁷ Currently not all advanced AI models can be run on device, although the capabilities of on-device or edge AI are rapidly improving.²⁸ Thus, AI agents are likely to employ hybrid solutions, where some tasks are performed by using AI on-device and others require *access to AI models deployed in the cloud*. Due to partnerships or vertical integration of developers of AI models with cloud providers, the cloud provider may often depend on which AI model is used, but this is not inevitable. Further, developers of AI agents also require computing resources (in the cloud) to build agentic frameworks and/or train AI models. However, training AI models, especially foundation models, involves a complex value chain and resources of its own – one that relies heavily on existing gatekeepers in the digital economy. The CMA highlights that current gatekeepers under the DMA are active (sometimes only through partnerships) in all parts of the value chain, ranging from access to computing resources and data in order to train models, over development of foundation models, to model deployment in specific application areas such as search, mobile ecosystems or productivity software.²⁹ This deep involvement of current gatekeepers has led commentators to conclude that there is "no AI without Big Tech" and that especially new startups have a high dependency on existing digital gatekeepers.³⁰

As noted above, agentic AI systems could rely on models provided by third-parties, and competition for agentic AI could therefore revolve around the agentic frameworks, rather than foundation models. However, as also noted, it seems that especially the providers of foundation models are currently pursuing the development and advancement of AI agents. Thus, while we have described the key components of an agentic AI system in a modular fashion, **vertical integration** between several of these key components is likely in practice. Vertical integration is likely to occur between the agentic framework and AI (foundation) models, between agentic framework and OS and/or hardware providers, or both. As the market is still evolving, the extent of vertical integration in the long-run is yet to be determined. The experience of mobile ecosystems has shown that full-stack vertical

²⁵ We borrow here from Anthropic's description of the key components of "Computer Use", see <https://docs.anthropic.com/en/docs/agents-and-tools/computer-use>.

²⁶ The Model Context Protocol (MCP) is one example of an emerging standard, see <https://www.anthropic.com/news/model-context-protocol>.

²⁷ See <https://www.qualcomm.com/news/onq/2024/06/whats-next-in-on-device-generative-ai>.

²⁸ Durga Malladi and Jerry Chang, 'AI Disruption is driving innovation in on-device inference' (Qualcomm White Paper, 2025).

²⁹ Competition & Markets Authority, Initial Report and Update Paper (n 8).

³⁰ Amba Kak, Sarah Myers West and Meredith Whittaker, 'Make no mistake—AI is owned by Big Tech' (MIT Technology Review, 5 December 2023) <https://www.technologyreview.com/2023/12/05/1084393/make-no-mistake-ai-is-owned-by-big-tech/>.



Is the DMA Ready for Agentic AI?

integration, ranging from chips/hardware over OS and software, like in the case of Apple, is feasible, despite complex value chains.



3. Potential Competition Concerns

The way in which AI agents are developed and deployed can raise competition concerns. Competition authorities studying (generative) AI have mostly focused on the development of foundation models – not on the specific issues raised by AI agents.³¹ This is not a problem, as most potential issues are known concerns manifesting themselves in a new setting. One can divide these concerns in two broad categories. First, a firm may promote its own AI agent and/or exclude that of competitors. Doing so requires market power either upstream (over one of AI agents' key inputs) or downstream (over important distribution channels). Second, a firm can use the AI agent itself to engage in anticompetitive conduct, e.g. by using it to self-preference products further downstream (e.g. specific apps). In order to do so, the firm operating the AI agent must have substantial power in the market in question – a market that may be broad (e.g. AI agents) but could also be narrower when consumers are locked in. The latter situation may arise when AI agents are integrated in another product (e.g. mobile devices) or if the costs of switching between agents is high (e.g. because their years of built-up memory are highly relevant). Let us look at these two categories of potential issues – foreclosure of an AI agent (Section 3.1) and foreclosure by an AI agent (Section 3.2) – in turn.

3.1. Foreclosure of an AI Agent

Foreclosure of an AI agent can proceed via control over key inputs or important distribution channels.

Input foreclosure

An AI agent has a number of essential building blocks (see above, Section 2). It is built on a foundation models that powers it. That foundation model is, in turn, trained by an expert workforce on a large dataset and with significant computing power (provided by chips directly or indirectly via a cloud service).³² Apart from the foundation model, the effective operation of an AI agent requires retrieving data from the environment (e.g. by using the camera or viewing your screen) and from other services (e.g. your calendar, e-mails, etc.) as well as computing power for inference, which can be provided on-device or through the cloud. The essential nature of each of these inputs depends on the evolution of the market but when an input does become essential, its owner can use it to foreclose other AI agent developers.

AI agent developers currently have a choice of leading **foundation models**. While developing one's own foundation model is an expensive undertaking, developers can access those of others either by relying on an open-source model (e.g. Meta's Llama)³³ or by relying on a more closed model that is made available through API calls (e.g. OpenAI's ChatGPT).³⁴ There is a risk that model providers cut off or degrade such access, e.g. as part of an 'open early, closed late' strategy.³⁵ The choice of AI model could also be limited when certain AI chips, embedded on devices, only support certain AI models.

³¹ For example, in Competition & Markets Authority, Initial Report (n 8), AI agents do not feature at all; in Competition & Markets Authority, Update Paper (n 8), they are mentioned in a footnote (n 25).

³² Competition & Markets Authority, Update Paper (n 8), paras 2.21–2.37.

³³ Even though Llama is not open-source according to (all) definitions, in particular because it does make its training data available, see <https://opensource.org/ai>.

³⁴ <https://openai.com/index/new-tools-for-building-agents/>.

³⁵ See Friso Bostoen and Nicolas Petit, 'Antitrust rules and remedies against platforms' treacherous turns' (2024) 30 European Law Journal 629.



The alternative to relying on a third-party foundation model, i.e. for agent providers to develop their own, may be difficult if – even abstracting from the high cost – the key inputs are not available.³⁶ **Compute**, both through chips³⁷ or cloud services³⁸, is concentrated but remains accessible. The availability of **training data** may be a larger stumbling block. While models are primarily trained on data crawled from the open web, websites can – and increasingly do – block crawlers.³⁹ Moreover, as the web has been crawled, training a competitive foundation model may increasingly require access to private datasets. Some firms developing AI agents already own such datasets (e.g. Google’s YouTube) and are likely unwilling to share them. An alternative is to license data from willing providers (e.g. social networks such as Reddit).⁴⁰ However, if licensing agreements become widespread and are conditioned on exclusivity with a single developer, other developers may not be able to get the qualitative data required to train their foundation model. In addition to compute and training data, specialized talent is also limited and hence expensive.⁴¹

Once powered by a foundation model, an AI agent needs a **carrier**. Unless agent-specific devices (e.g. Rabbit’s R1)⁴² break through, that carrier is likely to be the user’s desktop or mobile device. Access to a carrier straddles the line between input and distribution: a carrier is necessary to be put in front of users (i.e. to be distributed, see further under ‘distribution foreclosure’), but an agent developer also needs access to its specific hardware and software features (i.e. input, see the next paragraph). Arrangements between agent developers and carriers may be complicated if the device manufacturer or OS developer either has its own agent or has an (exclusive) partnership with another agent developer. In those cases, it may not want to give (the required degree of) access to the independent agent developer.

Even if the independent AI agent gets to “live” on the device/in the OS in question, it needs access to **hardware and software features** controlled by the device/OS provider and by other parties (e.g. developers of specific apps). Say a user asks their agent to buy the sweater that is being advertised on a billboard and to have it delivered when they are home. The agent can only seamlessly do so when it gets access to (i) the mobile device’s camera to snap the sweater; (ii) the user’s shopping app or browser to order the sweater; (iii) the user’s mobile payment mechanism to pay for the sweater; and (iv) the calendar app to check when the user is available for delivery. The incentives of each third party to grant interoperability are complex, even when they do not compete against the AI agent provider in question.⁴³ But wide access hardware, software *and* a user’s private information is crucial for an agent to effectively carry out tasks. Competition law sets a high bar for access to hardware and software (i.e. interoperability), though it has shifted somewhat in the *Android Auto* judgment.⁴⁴ If the owner of the hardware/software has reserved it for its own use, the stringent refusal to supply test of

³⁶ On the (continued) availability of key inputs, see also Meyers and Bourreau (n 10).

³⁷ The most high-performance chips are provided by a single firm, Nvidia, which is being scrutinized by competition authorities, see, e.g., <https://www.reuters.com/technology/french-competition-authority-confirms-investigation-into-nvidia-2024-07-15/>.

³⁸ The cloud market oligopolistic, with the majority in the hands of Amazon, Microsoft and Google, see Ofcom, ‘Cloud services market study’ (Final Report, October 2023).

³⁹ <https://reutersinstitute.politics.ox.ac.uk/how-many-news-websites-block-ai-crawlers>.

⁴⁰ <https://www.reuters.com/technology/reddit-ai-content-licensing-deal-with-google-sources-say-2024-02-22/>.

⁴¹ <https://www.wsj.com/tech/ai/meta-ai-recruiting-mark-zuckerberg-5c231f75>.

⁴² <https://www.rabbit.tech/>.

⁴³ App-based services, for example, are at risk of being commoditized (as in: users ask agents to simply do *something* – not to do so *via a specific app*). This has been called the ‘DoorDash problem’, see <https://www.theinformation.com/articles/openai-agents-pose-risks-doordash-consumer-apps>.

⁴⁴ Case C-233/23, Alphabet Inc. v Autorità Garante della Concorrenza e del Mercato (‘Android Auto’), EU:C:2025:110. Earlier, see Microsoft (Case COMP/C-3/37.792) Commission Decision of 21 April 2004 (on interoperability with Windows OS) and *Apple – Mobile Payments* (Case AT.40452) Commission Decision of 11 July 2024 (on interoperability with the iPhone’s near-field communication [NFC] chip).



indispensability (see below, Section 3.2) applies; if that hardware/software is developed with a view to enabling third parties (e.g. is an open platform), a lower standard requiring simple anticompetitive effects applies.

Distribution foreclosure

As noted above, an AI agent needs a carrier. This can be a device/OS, but it can also be another service, e.g. a search engine or social network.⁴⁵ Complications can arise if the provider of the carrier provides its own agent or has a preferred agent (e.g. due to an exclusive/revenue-sharing arrangement with another agent developer). In that case, the carrier may want to give prominent placement, via pre-installation or increased visibility, to its (preferred) agent.⁴⁶ By the same token, the carrier may want to make it more difficult for independent agent providers to reach their users. This can result in scenarios of distribution foreclosure, which have given rise to several antitrust infringement decisions in the digital economy over the past years. Let us take inspiration from them to get more specific on potential competition concerns.

First, an AI agent could be **pre-installed** (and made the **default**) on a popular device/OS, providing it with wide distribution. In previous cases, the pre-installation by an OS provider of their own app has been found abusive. Think, for example, of the pre-installation of Windows Media Player and Internet Explorer on Windows OS in, respectively, *Microsoft I* and *Microsoft II*.⁴⁷ In *Google Android*, the issue was also one of pre-installation, though the European Commission saw the tie in question differently: not between the OS (Android) and a specific app (e.g. Google Search) but rather between different apps (including Google Play, Chrome and Search).⁴⁸ In that same case, another abuse was found in the agreements between Google and device manufacturers that provided the latter with a share of the revenue if they (exclusively) pre-installed Google's services, in particular its search engine.⁴⁹

Second, an AI agent could be **integrated** into a popular service, which may also guarantee its broad availability. The key case on integrating one service into another is *Google Shopping*, which turned on the integration of Google's comparison shopping service in its search engine.⁵⁰ In particular, Google Shopping was given prominence by being shown on top of the search results with enhanced graphical features. *Facebook Marketplace* concerned an even more relevant scenario, in which Meta integrated Facebook Marketplace into its social network.⁵¹ Legally, Google's conduct was qualified as 'favouring' (or 'self-preferencing') while Meta's conduct was conceived as a **tie**.

Whatever the exact legal qualification, the pre-installation or integration of a specific app/service can become anticompetitive when it provides said app/service with a significant competitive advantage

⁴⁵ Meta, for example, integrated its AI assistant into its messaging apps, see <https://about.fb.com/news/2025/03/europe-meet-your-newest-assistant-meta-ai/>. To be fully functional in the real-world, as in the above example, those apps still need access to a device's hardware and software features.

⁴⁶ Less prominent placement is a potential method of foreclosure for platforms that are (partially) open, such as OSs and search engines; for more closed platforms, like messaging, third-party AI agents likely do not get access at all.

⁴⁷ *Microsoft* (Case COMP/C-3/37.792) Commission Decision of 21 April 2004; *Microsoft (tying)* (Case COMP/C-3/39.530) Commission Decision of 16 December 2009.

⁴⁸ *Google Android* (Case AT.40099) Commission Decision of 18 July 2018. Different from the Microsoft cases, the tie here was not between an OS and the apps in question, but rather between the apps themselves (e.g., Google Play and Google Search).

⁴⁹ Note that the revenue-sharing limb of the decision was annulled on appeal, though the issue was evidentiary, see Case T-604/18, *Google and Alphabet v Commission (Google Android)*, EU:T:2022:541. In the U.S., see similarly *United States et al. v. Google LLC*, 8 August 2024 (D.D.C. 2024).

⁵⁰ *Google Search (Shopping)* (Case AT.39740) Commission Decision of 27 June 2017.

⁵¹ *Facebook Marketplace* (Case AT.40684) Commission Decision of 14 November 2024.



that cannot be offset by relying on other distribution channels.⁵² There is thus always a question of alternatives (other devices/OSs/services) through which the independent agent developer can reach users. If the dominant firm's distribution method makes its (preferred) agent ubiquitous, the threshold can be met. In addition, if the dominant provider of the device/OS/service in question not only gives prominence to its (preferred) agent but also actively excludes that of competitors (e.g. makes the competing agent unfindable), anticompetitive conduct becomes easier to recognise.

3.2. Foreclosure by an AI Agent

Foreclosure by an AI agent is possible because such agents make recommendations to the user and can even take decisions on behalf of the user. For example, in carrying out tasks for their users, they constantly make recommendations ("you're asking me to buy you a new Aran jumper; how do you like this option from Aran Sweater Market?"). As they become more autonomous and learn their users' preferences, they may make more of the decisions independently, which cuts out the recommendation stage altogether. This is likely to happen first for low-cost decisions, such as where to order the Margherita pizza the user asked for (via DoorDash or the Domino's website?),⁵³ but may gradually expand to more significant decisions.

AI agents' recommendation and decision power can be harnessed to prop up its developer's other services through **leveraging and tying**. If the agent developer also happens to run a marketplace or food-delivery service, it may send requests for sweater and pizza purchases in that direction. This was a concern that scholars already voiced with regard to voice assistants, which were to become "gatekeepers for consumption" that could "nudge" users towards their preferred (vertically integrated) services.⁵⁴ While such anticompetitive concerns were not borne out for voice assistants, they seem more realistic when it comes to AI agents.

A more accurate analogy, in terms of potential anticompetitive effects, may not be with a voice assistant but rather with a search engine, i.e. a service that is central to how users search for (and then act on) information. The precedent here is the aforementioned *Google Search* decision, in which the European Commission found Google's channelling of demand towards its own specialised shopping service abusive. The determinative question, in such **self-preferencing** cases, is whether the search engine/AI agent in question is an important source of traffic for the downstream service and whether the diverted traffic cannot effectively be replaced.⁵⁵

There is, however, an important difference between AI agents and search engines or other platforms that currently channel demand such as online marketplaces.⁵⁶ When a search engine or marketplace puts its own product on top of the results and competing ones further down, the self-preferencing is somewhat transparent. Users may still be under the (false) impression that the platform's own product is in first place because it is the best product, but at least they can see that the platform is pushing that product. When an AI agent makes decision for its user, the selection process may happen

⁵² See *ibid*, para 752 and the references there.

⁵³ This is the origin of the 'DoorDash problem', see above (n 44).

⁵⁴ Victoriia Noskova, 'Virtual assistants as gatekeepers for consumption? – how information intermediaries shape competition' (2023) 19 European Competition Journal 30 and Viktorija Morozovaite, 'The future of anticompetitive self-preferencing: analysis of hypernudging by voice assistants under article 102 TFEU' (2023) 19 European Competition Journal 410.

⁵⁵ *Google Search (Shopping)* (Case AT.39740) Commission Decision of 27 June 2017, Sections 7.2.2–7.2.4.

⁵⁶ Also noted by Hunt, Chissell and Mawar (n 5), 27.



entirely “behind the screens” (unless the user prompts the agent to explain its actions). The self-preferencing can thus go unnoticed and may thus be more effective as it cannot be countered by users not clicking the first result; they would have to give specific instructions to use this or that service, but many users may leave that selection work to the agent. The risk of autonomous decisions that benefit the agent’s provider but not the user is especially large when users do not pay for the agent, whose provider thus has to be remunerated by other parties (e.g. those to which it channels demand). The effects of self-preferencing may be greater when an AI agent is pre-installed or integrated in a device/service, in which case users – even when aware of bias – have a harder time “escaping” it.

Additional competitive concerns may arise when a “winner takes all” dynamic asserts itself and a single AI agent becomes *incontournable*. Such a dynamic may be spurred by indirect network effects: if third-party apps/services need to be reconfigured to be compatible with an AI agent, developers may only want to do so for a few agents, which can have a concentrating effect.⁵⁷ If the market tips towards a certain AI agent and its developer then **refuses to deal** with, for example, an OS provider that would like to integrate the agent, this can harm competition. An incentive to refuse in such scenarios only exists if the agent developer stands to gain, for example because its own OS is set to win market share as a consequence of the refusal. In such situations, the refusal to supply doctrine may apply.⁵⁸ It sets a high bar: the AI agent must be indispensable (meaning there are no actual or potential substitutes) and the refusal must eliminate all competition (on the part of the firm requesting access).

Finally, as users rely more on AI agents, and the agents learn its users’ preferences more and more over time, this bears the risk of **data-induced lock-in effects**. Users may be reluctant to switching to a provider of a rival agent if they cannot take their data and interaction history with them. Data-enabled learning can occur both within a user, but also across users, and can constitute a significant source of competitive advantage over other providers that do not have access to similar user data at scale.⁵⁹ Thus, data-driven network effects may further spur the “winner takes all” dynamic, and, especially if coupled with foreclosure of data access by the dominant AI agent provider likely further hampers rival agent providers’ ability to compete.

⁵⁷ Japan Fair Trade Commission, ‘Generative AI and Competition’ (Discussion Paper, October 2024), 11. Apple, for example, lets developers make their apps “discoverable” to AI agents with App Intents, see <https://developer.apple.com/documentation/appintents/>.

⁵⁸ See Joined cases 6 and 7-73 *Commercial Solvents v. Commission*, EU:C:1974:18; Joined cases C-241/91 P and C-242/91 P *RTE and ITP v. Commission* (‘*Magill*’), EU:C:1995:98; Case C-7/97 *Oscar Bronner v. Mediaprint*, EU:C:1998:569.

⁵⁹ See, e.g., Andrei Hagiu and Julian Wright, ‘Data-enabled learning, network effects, and competitive advantage’ (2023) 54 *The RAND Journal of Economics* 638; Krämer and Schnurr (n 22).



4. Is Agentic AI Within the DMA's Scope?

As the AI agents market is only nascent, the exact risk and magnitude of any competition concerns remains uncertain. In the previous section, we have sketched potential concerns based on the experience in other digital markets, accounting for the specific features of AI agents. In case anticompetitive effects do manifest themselves, we gave a first indication of the extent to which competition law applies. However, in view of competition law's limitations in digital markets, both substantive (the difficulty of defining markets, establishing dominance, proving abuse) and institutional (the long duration of proceedings and ineffectiveness of remedial processes), EU legislators adopted the DMA. Insofar as the DMA applies, it becomes – due to its procedural efficiency – the European Commission's first choice to tackle any competitive issues that may arise. In this section, we therefore ask whether agentic AI falls within the DMA's scope.

The DMA applies to gatekeepers, i.e. providers of CPSs that have been designated with gatekeeper status.⁶⁰ Thus, in order to be regulated under the DMA, a service must first meet the definition of at least one of the ten CPSs. Against the backdrop of preceding description of AI agents, we can immediately exclude “online social networking services”, “video-sharing platform services”, “number-independent interpersonal communications services” and “online advertising services” as conceivably fitting categories for the designation of AI agents. This leaves us with the following options for CPS categories that may lend themselves for interpretation to the context of agentic AI:

- online intermediation services;
- online search engines;
- OSs;
- web browsers;
- virtual assistants;
- cloud computing services.

We start by discussing whether these existing CPS categories accommodate AI agents (Section 4.1). Then, we consider the potential gatekeeper status of AI agents qualified as a CPS (Section 4.2), before providing a final analysis of the DMA's scope with regard to AI agents (Section 4.3).

4.1. Designation as Core Platform Service

As there are six CPSs which AI agents *may* qualify as, a preliminary question is whether a service needs to be designated as a single CPS or whether a combinatorial designation is possible. When setting thresholds for designation, the DMA does refer to CPSs in the singular. The issue came up in the designation of TikTok. ByteDance, TikTok's owner, argued that it should be designated as a video-sharing service; instead, the European Commission designated it as a social network.⁶¹ There were valid arguments for either position: video-sharing is TikTok's primary purpose and core function, while

⁶⁰ Regulation (EU) 2022/1925 of the European Parliament and of the Council of 14 September 2022 on contestable and fair markets in the digital sector and amending Directives (EU) 2019/1937 and (EU) 2020/1828 (Digital Markets Act) [2022] OJ L265/1 (hereafter: Digital Markets Act), arts 2(1) and 3.

⁶¹ *ByteDance – online social networking services* (Case DMA.100040) Commission Decision of 5 September 2023.



its breadth of other features brings it closer to a social network. Neither ByteDance nor the Commission explored a dual designation and, as ByteDance did not dispute its qualification on appeal, we did not hear the General Court's opinion.⁶² Hence, also in view of the practical difficulty combinatorial designation would cause (e.g. when counting users), it seems safe to assume a service can currently only be designated in one CPS category. This makes the finding the *right* category all the more important, especially as the obligations also vary per category (see below, Section 5).

Online Intermediation Service

Online intermediation services are defined with reference to the P2B Regulation as (i) information society services⁶³ (ii) that allow business users to offer goods or services to consumers with a view to facilitating direct transactions between those business users and consumers, and (iii) that are provided to business users on the basis of contractual relationships between the provider of those services and business users which offer goods or services to consumers.⁶⁴

The definition emphasises intermediation between consumers, on one side, and business users (offering products or services to consumers via the online intermediation service), on the other. While an AI agent may be employed as a tool to discover and transact with businesses (e.g. an AI shopping agent), typically there is no contractual relationship between the business user (e.g. third-party shops) and the AI agent. Rather, the agent connects to applications or, indeed, online intermediation services (e.g. online marketplaces); alternatively, it surfs the web to complete the purchase.⁶⁵ The applications or websites visited by the agent may have business users, but the agent itself does not (directly).

The provider of the AI agent *could* have contractual relations with some of the business users that it may interact with, but the nature of the AI agent is such that – following a user's request – it would autonomously search for relevant matches and perform actions independent of pre-existing contractual relations with the parties that it interacts with. The presence of contractual relations with business users is not a defining feature of AI agents.

In conclusion, as AI agents do not directly host business users, nor tend to have contractual relations with them, they generally do not qualify as an online intermediation service in the sense of the DMA.

Online Search Engine

The DMA defines online search engines also with reference to the P2B regulation as “digital service[s] that allow ... users to input queries in order to perform searches of, in principle, all websites ... on the basis of a query on any subject in the form of a keyword, voice request, phrase or other input, and returns results in any format in which information related to the requested content can be found”.⁶⁶

⁶² Case T-1077/23, *Bytedance v Commission*, EU:T:2024:478, para 29.

⁶³ Defined by reference to Directive (EU) 2015/1535 of the European Parliament and of the Council of 9 September 2015 laying down a procedure for the provision of information in the field of technical regulations and of rules on Information Society services [2015] OJ L241/1, art 1(1)(b) (“any service normally provided for remuneration, at a distance, by electronic means and at the individual request of a recipient of services”).

⁶⁴ Regulation (EU) 2019/1150 of the European Parliament and of the Council of 20 June 2019 on promoting fairness and transparency for business users of online intermediation services [2019] OJ L186/57 (hereafter: P2B Regulation), art 2(2).

⁶⁵ This is the model of Amazon's “Buy for Me” features, which makes use of “agentic AI”, see <https://www.aboutamazon.com/news/retail/amazon-shopping-app-buy-for-me-brands>.

⁶⁶ P2B Regulation, art 2(5), referenced in Digital Markets Act, art 2(6).



This definition contains some of the core features of an AI agent. Indeed, an integral part of AI agents is that they are able to autonomously retrieve external information, such as those provided by, in principle, all websites.⁶⁷ This is also becoming a prevalent feature of pre-trained AI foundation models. In order to combat hallucination and to overcome the knowledge cut-off inherent to the training data, foundation models will also resort to new information collected from the web as needed. Web search has therefore become an integral part of modern LLM-based AI systems. However, the functionality of AI agents is, of course, not limited to searching the web, and goes beyond information retrieval. We will return to this point below.

One could also argue that AI agents, particularly those embedded in personal devices, perform their actions in response to a “query on any subject”, i.e. a user request. The definition of search engines allows for a wide range of possibilities in which this request can be formulated, including text, voice or “other input”. Thereby, “other input” can be interpreted as a catch-all term that possibly opens the door for all additional conceivable options in which a human may interact with technology, such as AI agents, including, but not limited to gestures, facial expressions or even through neural interfaces.

Finally, the format in which the search engine must provide the results in response to a user request is also defined in a quite open-ended fashion. Indeed, it can be in “any format”. Therefore, the form in which the AI agent responds and delivers “results” to the users’ request does not disqualify it from being a search engine.

For example, a user that formulates a request to the AI agent to “buy the cheapest tickets for the next concert of Shania Twain”, which requires the AI agent to “search” for cheap tickets on, in principle, all websites (ticket sellers, private listings, online marketplaces, online communities, etc.) and returns results in form of a valid (electronic) ticket, may awkwardly fit the definition of “search engine” under the DMA.

However, as mentioned above, the nature of an AI agent is not limited to search functionality. It also caters to requests that do not necessarily involve the search of “in principle, all websites”. Further, albeit AI agents are likely to encompass a search function, and in turn all major search engines tightly integrate with AI, the integration between a search engine and the AI agent can vary. The few providers of AI agents that have their own search engine can rely on it, providing for a tight integration (e.g. Bing AI, Google Gemini). Other providers have to rely on one of the main search engines (e.g. OpenAI’s ChatGPT taps into Bing). In either case, the question is whether the AI agent *is* a search engine or simply *makes use of* a search engine.

There is a fine line between the two, which is getting even more blurred as AI agent providers launch ever more search-like products (e.g. OpenAI’s SearchGPT).⁶⁸ There are therefore valid arguments to qualify and disqualify AI agents as search engines. As long as standalone search engines⁶⁹ exist and are distinguishable from agentic AI systems, one may argue that AI agents are indeed by nature different

⁶⁷ A process called “retrieval augmented generation” (RAG), see <https://aws.amazon.com/what-is/retrieval-augmented-generation/>. Some AI agents have this option toggled on by default (Perplexity), others require the user to select it (ChatGPT).

⁶⁸ See <https://openai.com/index/searchgpt-prototype/>.

⁶⁹ It is worth mentioning that a search engine is not one monolithic entity either, but comprises several key components, which can, in principle, be separated. This includes, at a high level of abstraction, the web index (i.e., the directory of all websites), the ranking algorithm (which translates queries into a ranking of websites), and the user interface. Indeed, several smaller search engines, such as DuckDuckGo or Ecosia provide their own user interface, and may process user data differently than their larger rivals, but yet rely to a large extent on syndication agreements with Microsoft’s Bing (accessing, at least in part, their web index and ranking algorithm) in order to provide their search results.



to a search engine and simply make use of a search engine. However, given the rapid evolution and enhancement of search engines in response to AI, in a not-so-distant future, standalone search engines, as we know them, may no longer exist or at least no longer play a significant role. Instead, search is likely to become a commodity fully embedded in AI-based consumer-facing services, such as AI assistants or indeed AI agents. In this case, search – as envisioned by the DMA – could only be performed through these AI-services, and thus AI agents may well be considered (at least in part) a search engine.

Operating System

In the DMA, an OS is defined as “a system software that controls the basic functions of the hardware or software and enables software applications to run on it”.⁷⁰ At present, such a qualification seems far-fetched: AI agents are, at most, an app or website accessed via an OS – they are not themselves OSs. Looking to the future, however, that might change. First, as noted earlier, for an AI agent to operate effectively, it needs access to a device’s hardware and software. Over time, it may not just get access but actual control. Already, firms are deeply embedding these agents in their OSs. Apple Intelligence, for example, is said to be “integrated into the core of your iPhone, iPad, and Mac”.⁷¹ On the other side, developers can use Apple’s App Intents framework to “deeply integrate” their app’s functions and content with Apple Intelligence.⁷² As a result, Apple Intelligence can call upon the app for specific tasks without the user having to navigate there. However, while the AI is invoking the apps, the apps are not yet “running on” the AI agent.

In the medium-to long-term, AI agents may not just run apps but also take over control over the device’s hardware and software from the OS – in effect, become the OS. Commentators are already predicting “the rise of AI OS” and “AIs in your OS”.⁷³ A (for now academic) model is found in MemGPT, which draws inspiration from the hierarchical memory systems in traditional OSs to increase LLMs context windows.⁷⁴ While such developments are likely to bring AI agents closer to OSs over time, they are unlikely to qualify as OSs in the sense of the DMA in the short-term.

Web Browser

The DMA defines a web browser as “a software application that enables end users to access and interact with web content hosted on servers that are connected to networks such as the Internet, including standalone web browsers as well as web browsers integrated or embedded in software or similar”.⁷⁵

Reminiscent of our discussion of AI agents as search engines, AI agents may well embed a web browser as an integral component. The functionality of the current prototypes OpenAI Operator, Claude Computer Use and Google’s Project Mariner all heavily rely on a web browser to perform their tasks. Thereby the definition of a web browser in the DMA is very broad, similar to that of search engines, encompassing virtually any software that allows to connect to “web content”, whether it be on the (public) Internet or any other servers connected to (private) networks. Notably, the definition is

⁷⁰ Digital Markets Act, art. 2(10).

⁷¹ <https://www.apple.com/apple-intelligence/>.

⁷² <https://developer.apple.com/documentation/appintents/>.

⁷³ <https://www.walturn.com/insights/the-rise-of-ai-os>, <https://medium.com/@michaeldain/ais-in-your-os-0bbc87821210>.

⁷⁴ Charles Packer et al., ‘MemGPT: Towards LLMs as Operating Systems’ (2024) arXiv:2310.08560.

⁷⁵ Digital Markets Act, art 2(11).



explicitly not limited to servers connected to the Internet. This is a bit odd in conjunction with the use of the word “web content”, as “web” is commonly understood as an abbreviation of “World Wide Web” (WWW), which commonly refers to a system of interconnected *public* webpages accessible through the *Internet*.⁷⁶ A more narrow definition of “web content” could mean content accessible through the Hypertext Transfer Protocol (HTTP), which can also be used in private networks. However, modern browsers are powerful pieces of software,⁷⁷ encompassing a rich set of features and functionalities not limited to HTTP, even if all browsers – as their core functionality – support HTTP.

Similar as in the case of search engines, the functionality of AI agents clearly goes beyond what is commonly considered a web browser today. Again, the central question is therefore whether an AI agent *is* a web browser or just *makes use of* a web browser.

Looking ahead, it is likely that AI agents will fundamentally challenge the notion of standalone browsers as we know them. Searching and browsing the web will be done for us, rather than by us, as in the case of ChatGPT Operator and the like. To facilitate this task, firms like OpenAI may even develop their own version of a “browser”.⁷⁸ In any case, the distinction between AI agents and browser will increasingly become blurred, just like in the case of search engines.

Virtual Assistant

A CPS category particularly reminiscent of AI agents is that of “virtual assistant”.⁷⁹ The DMA defines a virtual assistant as “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”.⁸⁰ Whether AI agents are captured by the DMA as virtual assistants depends in part on the relative weight one gives to the legislative text versus the underlying intent at the time of adoption.

Starting with the text, the definition of virtual assistant is – as that of search engines and browsers – very broad. In its plain meaning, the definition would capture AI agents, whose essence is to process tasks or questions based on a variety of cues (voice, writing, gestures, etc.). Semantically, some make a distinction between assistants and agents. IBM, for example, puts it as follows: “AI assistants are reactive, performing tasks at your request. AI agents are proactive, working autonomously to achieve a specific goal by any means at their disposal.”⁸¹ This semantic distinction is unlikely to matter, however, as qualification is based on substance – not labels. Moreover, the distinction obscures that assistants and agents are situated on a spectrum, with current LLM-based services moving from the assistant to the agent side. Based on the substance of the definition, AI assistants and the agents they are evolving into may thus qualify as virtual assistants. In the more distant future, as agents reach a high degree of autonomy, the definition may become less fitting as there may no longer be (explicit)

⁷⁶ See, e.g., https://en.wikipedia.org/wiki/World_Wide_Web.

⁷⁷ Similar as in the case of search engines, browsers are comprised of several distinguishable and separable components, including, at a high level of abstraction, the rendering engine (where websites are rendered, such as Blink for Chrome, Edge and Opera; WebKit for Safari; Gecko for Firefox), the user interface and the browser engine (which handles interactions between the user interface and the rendering engine), see <https://dev.to/itswillt/the-components-of-a-browser-23mn>.

⁷⁸ <https://www.linkedin.com/pulse/why-openai-has-build-browser-jonathan-milne-qwssc/>.

⁷⁹ Digital Markets Act, art 2(2)(h).

⁸⁰ Digital Markets Act, art 2(12).

⁸¹ See <https://www.ibm.com/think/topics/ai-agents-vs-ai-assistants>.



“demands, tasks or questions” (prompts) given by the user; the AI agent will proactively take action based on goals previously set by the user or even just inferred by the agent.⁸²

The legislative history somewhat complicates this picture. Virtual assistants were not included in the original DMA proposal⁸³ but were later added by lawmakers.⁸⁴ At the time, the European Commission was carrying out a sector inquiry into the consumer Internet of Things (IoT), with a focus on voice assistants such as Amazon’s Alexa.⁸⁵ While there were competitive concerns, e.g. around interoperability, the Commission did not launch a subsequent antitrust investigation. Concerns did find their way into the DMA’s legislative process,⁸⁶ where the European Parliament picked them up. Lawmakers included a definition of virtual assistants that is much broader than the voice assistants scrutinised at the time of the DMA’s adoption but, since this happened in trilogue, the reasons for doing so are not public. Nevertheless, there are some signs in the DMA itself that lawmakers had traditional voice assistants – not AI agents – in mind during the drafting process. First, the “virtual assistant” definition refers to the control of connected physical devices, which is a typical feature of Alexa-type voice assistants.⁸⁷ Second, the DMA defines business users of a virtual assistant as “developers who offered at least one virtual assistant software application or a functionality to make an existing software application accessible through the virtual assistant”.⁸⁸ This seems to hint at the “skills” developed for specific voice assistants like Alexa.⁸⁹ The idea of AI agents, by contrast, is that they can autonomously access various apps, which does not necessarily require apps to be developed specifically for the AI agent. We are, however, seeing the first signs of agent-specific API frameworks, and of partnerships between agent and app developers.⁹⁰

In conclusion, the DMA’s “virtual assistant” category provides a rather close fit for AI agents, at least in the near to mid-term future. The definition certainly accommodates them, even if it is unclear what the legislative intent was when adding this specific CPS to the DMA. In the longer-term, and in view of policy considerations such as legal and investment certainty, a legislative rather than interpretative route may be preferable. We return to this in our concluding analysis (Section 4.3).

Cloud Computing Service

The DMA defines cloud computing services in reference to EU Directive 2016/1148 as “a digital service that enables access to a scalable and elastic pool of shareable computing resources”.⁹¹ Even though

⁸² In case of goals set by the user, the “virtual assistant” definition could still apply; in case of goals inferred by the agent, that becomes difficult.

⁸³ See Proposal for a Regulation of the European Parliament and of the Council on contestable and fair markets in the digital sector (Digital Markets Act) COM/2020/842 final.

⁸⁴ European Parliament, ‘Digital Markets Act: ending unfair practices of big online platforms’ (press release, 23 November 2021) <https://www.europarl.europa.eu/news/en/press-room/20211118IPR17636/digital-markets-act-ending-unfair-practices-of-big-online-platforms>.

⁸⁵ The inquiry started in 2020 and was finalized in 2022, see European Commission, ‘Sector Inquiry into Consumer Internet of Things’ (Final Report) COM(2022)19 final.

⁸⁶ European Commission, Impact Assessment Report accompanying the Proposal for a Regulation of the European Parliament and of the Council on contestable and fair markets in the digital sector (Staff Working Document) SWD(2020)363 final, para 97.

⁸⁷ Similarly, Digital Markets Act, Annex, Section E, refers to “controlling a smart home device”.

⁸⁸ Ibid.

⁸⁹ <https://www.amazon.com/gp/help/customer/display.html?nodeId=G8QQCQL47RFVGAZC>. Note that some AI agents have plug-ins, which are similar to voice assistant “skills”. OpenAI had them but has meanwhile deprecated them, see <https://openai.com/index/chatgpt-plugins/>.

⁹⁰ For an example of an agent-specific API framework, see <https://developer.apple.com/documentation/appintents/>. As for partnerships, OpenAI is – for its Operator agent – partnering with DoorDash, Instacart, OpenTable, Priceline, StubHub, Thumbtack, Uber and others, see <https://openai.com/index/introducing-operator/>.

⁹¹ Digital Markets Act, art 2(13), by reference to Directive (EU) 2016/1148 of the European Parliament and of the Council concerning measures for a high common level of security of network and information systems across the Union [2016] OJ L194/1, art 4(19).



AI agents may in their functionality heavily rely on access to cloud computing services, their nature is not to “enable access to” sharable computing resources. Rather AI agents are users of cloud computing services, even though cloud computing services and AI agents may be provided by the same company. Even in view of future developments, the designation of AI agents as a cloud computing service does not seem warranted.

4.2. Determination of Gatekeeper Status

After settling on a fitting CPS category, the AI agent provider in question would need to achieve gatekeeper status in order to fall under the regulation of the DMA. Gatekeeper status under the DMA is conditioned on three qualitative criteria: (a) having a significant impact on the internal market; (b) serving as important gateway for business users to reach end-users; and (c) enjoying an entrenched and durable position.⁹² These criteria are presumed to be fulfilled when a firm meets certain quantitative criteria, respectively:

- (a) Achieving EU revenue of at least €7,5 billion in each of the last three financial years or having an average market capitalisation/fair market value of at least €75 billion in the last financial year, and providing the same CPS in at least three Member States;
- (b) Providing a CPS that has at least 45 million monthly active end-users and 10.000 yearly active business users in the EU in the last financial year; and
- (c) Meeting the threshold of (b) in each of the last three financial years.⁹³

Abstracting from the CPS qualification question discussed above (Section 4.1), these thresholds raise some questions when it comes to AI agents. Criterion (a), on **revenue/market cap** (or fair market value) remains straightforward. The interpretation difficulty is centred on criterion (b), regarding the **number of end- and business users**. Identifying users that have engaged with the AI agent in the past month should pose no difficulty, but identifying business users does.

According to the DMA, a business user is “any natural or legal person acting in a commercial or professional capacity using core platform services for the purpose of or in the course of providing goods or services to end users”.⁹⁴ This definition seems to imply that a business user must actively choose the CPS to conduct business. The specific definition of business users of virtual assistants, discussed already above (under “Virtual Assistant”) also appears to require active conduct, as it covers “unique developers who offered at least one virtual assistant software application or a functionality to make an existing software application accessible through the virtual assistant”.⁹⁵

In the context of AI agents, however, it is questionable whether the business user has actively chosen the AI agent to do business, or the other way around (see already above under “Online Intermediation Service”). Certainly, some business users are actively relying on AI agents via partnerships and/or API integration.⁹⁶ But “computer use” models seek out business users without their active cooperation.

⁹² Digital Markets Act, art 3(1).

⁹³ Digital Markets Act, art 3(2).

⁹⁴ Digital Markets Act, art 2(21).

⁹⁵ Digital Markets Act, Annex, Section E.

⁹⁶ See <https://developer.apple.com/documentation/appintents/> for APIs and <https://openai.com/index/introducing-operator/> for partnerships.



Business user identification thus poses a potential problem when designating AI agents as virtual assistants. One could get around this problem with a wider interpretation of business user, which considers any business called on by an AI agent to be a business user, given that this business is using a CPS “in the course of providing goods or services to end users”. This is the approach taken by the DMA for search engines, where a business user is anyone “with business websites *indexed by* the online search engine”.⁹⁷ Clearly, the search engine is the party doing the indexing, whereas the business websites are passively being indexed (though they can take active steps to resist such indexing).⁹⁸ It is therefore not unreasonable to analogously interpret business users of virtual assistants as encompassing also those that are chosen by the AI agent to provide goods or services to end-users, even if this requires a bit of a leap.

Criterion (c) sets an **entrenchment** threshold, i.e. boasting the above number of business and end-users *for three years*. There is thus a time-lag built into the DMA, based on the idea that – even when certain CPSs are seeing (momentarily) high user numbers – they need to show a degree of stability before being designated with gatekeeper status and subjected to regulatory obligations. Any new technology, including AI agents, will thus see the designation of gatekeepers at the earliest three years after its widespread adoption. There are, however, two exceptions to this rule.

First, the European Commission could designate an AI agent based on a **market investigation**, in which it establishes that the qualitative but not the quantitative gatekeeper criteria are met.⁹⁹ It would then still have to show that, based on the DMA’s criteria and barriers to entry, the provider is entrenched with respect to its AI agent. Given that AI agents are nascent and that the market is, for now, still heavily contested, this seems like a tall order.

Second, the Commission could designate an AI agent provider as “**emerging gatekeeper**”. *Current* entrenchment is not required in this case, rather, the question is whether the AI agent provider “will foreseeably enjoy [an entrenched and durable] position in the near future”.¹⁰⁰ This also requires a market investigation, and the Commission can only declare specific obligations applicable to the emerging gatekeeper.¹⁰¹ At the time of writing, no CPS provider has been designated as emerging gatekeeper, and no AI agent seems to fit that bill either. However, given the quick development of AI this may change very quickly.

4.3. Analysis

AI agents challenge the DMA’s CPS categorisation. Of the six CPSs that appeared to provide a plausible categorisation (i.e. could not immediately be excluded), there are – at present – three categories that more or less fit. Online search engines and web browsers have both been defined broadly and, as such, could be interpreted to cover certain AI agents. The connection lies in the fact that AI agents widely make use of the web for the retrieval of information and to carry out tasks. Nevertheless, it appears more correct to say that AI agents *rely* on search engines/browsers but do not *constitute*

⁹⁷ Digital Markets Act, Annex, Section E (own emphasis).

⁹⁸ See e.g., <https://developers.google.com/search/docs/crawling-indexing/block-indexing>.

⁹⁹ Digital Markets Act, arts 3(8) and 17(1)–(2). At the time of writing, there had been one such designation: *Apple – iPadOS* (Case DMA.100047) Commission Decision of 29 April 2024.

¹⁰⁰ Digital Markets Act, arts 3(8) and 17(4).

¹⁰¹ In particular, art 5(3)–(6) and art 6(4), (7), (9), (10) and (13), see Digital Markets Act, art 17(4).



them. However, that line becomes increasingly difficult to draw as AI agents integrate more and more search/browsing functionality.

A third CPS category, of virtual assistants, provides a cleaner fit. There are some questions about the legislative origins of this category but, given that this history is not transparent, it cannot be a significant source of authority, which leaves us with the text. The definition of virtual assistants is – like that of search engines and web browsers – broad. The DMA’s wide definitions suggest an attempt to keep the instrument reasonably open to the rapid evolution of digital markets, in other words, to keep it future-proof. Hence, the broad virtual assistant definition – despite textual remnants of the voice assistants that legislators had in mind during the DMA’s adoption – relatively comfortably accommodates AI agents.

This finding must be nuanced in two ways. First, AI agents are rapidly evolving. While the virtual assistant definition may cover the first generations of AI agents, that coverage can become tenuous over time. As AI agents attain ever more capabilities, autonomous action is likely to become more prevalent: agents may shift from processing user input/prompts (“demands, tasks or questions”) to anticipating user “needs”, at which point the virtual assistant definition becomes less fitting. Second, for several CPS categories including that of virtual assistant, the Annex’s business user definition leads to issues, as AI agents do not necessarily or actively have them, even though this may also change over time if partnerships and API integration become more established practice.

The alternative to dealing with these issues through interpretation is legislative change. Explicitly adding AI agents to the list of CPS (by changing or replacing the virtual assistant definition) would require the Commission to carry out a market investigation and to submit it, accompanied by a legislative proposal, to the European Parliament and the Council.¹⁰² This can be done as part of the DMA’s three-yearly evaluation, of which the first one must be carried out by May 2026.¹⁰³ Adapting the definition of business users, in particular those of virtual assistants, is possible via a delegated act and is therefore less legislatively demanding.¹⁰⁴

In the choice between the two routes to capture AI agents under the DMA – interpretation and legislative change – there are different policy considerations at play. The main ones are effectiveness and legal certainty – two values that are central to the DMA.¹⁰⁵ Interpretation provides for (immediate) effectiveness; legislative change comes with greater legal certainty, which in turns provides confidence to investors. In addition, legislative change confers greater legitimacy on the inclusion of AI agents within the DMA’s scope, which may be particularly important in view of (i) the scant attention given to AI in developing the DMA, and (ii) the (political) calls for more caution in creating regulatory burdens in the EU (from within the EU and from across the Atlantic).¹⁰⁶

The two routes are not mutually exclusive. If AI assistants gradually evolve into agents, and become entrenched and important gateways for business users to reach end-users, the Commission could designate them under the virtual assistant CPS category in the near to mid-term. This ensures

¹⁰² Digital Markets Act, art 19. Obligations can be updated via delegated act, see art 12; there is no equivalent procedure for updating (as different from adding) CPSs.

¹⁰³ Ibid, art 53.

¹⁰⁴ Ibid, arts 3(7) and 49.

¹⁰⁵ See, e.g., ibid, recital 77 on the need to ensure that the DMA both “constitutes an effective and holistic regulatory response” and “ensure[s] legal certainty as to the regulatory conditions”.

¹⁰⁶ On transatlantic pressures, see <https://www.politico.eu/article/grow-us-eu-rules-tech-donald-trump-pressure-brussels/>.



effectiveness and, given that the first generations of AI agents are close (enough) to the virtual assistants described in the DMA, legal certainty is not significantly imperilled. Before designation, however, it would be useful to adapt the business user definition via delegated act. If agents become ever more autonomous, which may be expected in the mid to long-term, legislative change would bring greater legal certainty and might even be needed to effectively capture AI agents. At that point, such change can be based on a clearer view of the trajectory of agentic AI, e.g. towards more system-level software. This two-step approach rhymes effectiveness with legal certainty.

Moving from CPS qualification to gatekeeper status, there are fewer difficulties on that front (apart from defining and consequently “counting” business users, discussed above). Designation is – by design – a matter of time, as the DMA does not intervene before significant adoption. Market investigations and the status of “emerging gatekeeper” providing a possibility to modestly speed up that timeline, but we see no immediate need to rely on them. Finally, should AI agents meet the designation thresholds while the market nevertheless remains contestable (e.g. due to a lack of end-user lock-in and business user dependence), rebuttals must be considered.¹⁰⁷

¹⁰⁷ Digital Markets Act, art 3(5).



5. Are the DMA's Obligations Meaningful in the Context of Agentic AI?

In the preceding section we have determined that AI agents may conceivably be designated – provided they meet the criteria for (emerging) gatekeeper status – as either search engine, web browser, or virtual assistant, although the virtual assistant category is seemingly the most fitting of these. Over time, AI agents may come to resemble OSs, and they any case need some level of access to OSs to function. However, as none of these CPS categories were designed with the specific case of AI agents in mind, it is interesting to see whether the specific obligations that are attached to these categories, individually or collectively, are yet meaningful in the context of mitigating the competition concerns associated with agentic AI (Section 3).

While the DMA's obligations are sometimes seen as generally applicable to gatekeepers, many of them apply only to a specific CPS (or set of CPSs). As pointed out above (Section 4), this means that CPS qualification must not only be formally correct but also functionally relevant (i.e., in a category where obligations are relevant to the service in question). To underline this importance, we discuss the relevant obligations – as they pertain to AI agents – based on the CPS to which they apply: OS (Section 5.1); OS, virtual assistant, and web browser (Section 5.2); search engine (Section 5.3); all CPSs (Section 5.4). We do not discuss obligations that apply only to CPS categories in which AI agents cannot conceivably be designated.¹⁰⁸

5.1. Obligations for OSs

For OSs, the main obligations centre on users' ability to (un)install.

Uninstallation

The DMA contains specific obligations that only apply to a designated OS, and not to any other CPS category. First this is **Article 6(3), paragraph 1**, DMA, which demands that “the gatekeeper shall allow and technically enable end users to easily un-install any software applications on the operating system of the gatekeeper” unless those software applications are “essential”, meaning that they cannot reasonably offered by a third-party provider on a standalone basis.

Provided that AI agents are *not* designated under the DMA as an OS, but the OS provider has pre-installed a (non-essential) AI agent, then this obligation (in conjunction with the interoperability requirements of Article 6(7) DMA discussed further below) ensures that an alternative third-party AI agent can be installed on the OS. The provision is therefore meaningful in the context of preventing the foreclosure of a competitively supplied AI agent by an OS provider (see above, Section 3.1). Note that it is not material for this provision whether the pre-installed software is from the gatekeeper or not. Thus, the provision should also apply to cases where an OS provider merely cooperates with the provider of an AI agent, e.g., in case of Apple cooperating with OpenAI for Apple Intelligence.

However, as agentic AI and the OS are likely to become more intertwined (see above, Section 4.1), a deep integration of the AI agent and the OS is required, so there will be debate around how “essential”

¹⁰⁸ E.g., Digital Markets Act, arts 5(9)–(10) and 6(8), which only apply to advertising services.



the pre-installed AI agent is, and whether it could indeed be offered on a standalone basis. The latest advancements in OS development, both on mobile and desktop devices, already integrate AI support to a degree that well exceeds the integration of other typical third-party applications. If the embedding of AI agents into the OS is indeed deemed “essential”, then this opens up the possibility to designate AI agents as an OS. In this case, there would be no uninstallation requirement according to Article 6(3), paragraph 1, but the gatekeeper providing the integrated OS/AI agent environment would yet need to provide equitable interoperability to third-party AI agent providers according to Article 6(7) DMA, as discussed below, since the OS-integrated AI agent is still a “software feature” provided by the gatekeeper. However, consumers may be more reluctant to installing a third-party AI agent in case the pre-installed AI agents cannot be un-installed – which was precisely the legislative intent behind the uninstallation requirement expressed in the first paragraph of Article 6(3).

Installation

For completeness, we mention that **Article 6(4) DMA** also entails a provision that requires “the gatekeeper [to] allow and technically enable the installation and effective use of third-party software applications ... using, or interoperating with, *its operating system* and allow those software applications ... to be accessed by means other than the *relevant core platform services* of that gatekeeper.”¹⁰⁹ Although the obligation clearly requires that the gatekeeper has an OS (whether designated as CPS or not), this is not the “relevant” CPS. This provision was introduced thinking about app stores being the relevant CPS, and it only seems to make sense in this specific context, unless, in the future, users would only be able to access third-party software applications through the use of an AI agent, which would then need to be designated in its own right (separate from the OS). This seems a far-fetched scenario from today’s vantage point.¹¹⁰

5.2. Obligations for OSs, Virtual Assistants, Web Browsers

In the DMA, virtual assistants are often put in a category different from “simple” applications.¹¹¹ In particular, they are lumped together with OSs and browsers, as the type of software that *other* application run on.¹¹² For system software, such as OSs, that function is clear. The inclusion of browsers in this category is understandable: the (US) *Microsoft* case was about how Microsoft saw a “middleware” threat in Netscape’s browser, which could become a platform for applications in its own right and thus make Windows obsolescent.¹¹³ The inclusion of virtual assistants in this category is – or was – not immediately obvious. With its rudimentary “skills”, for example, Amazon’s Alexa did not look likely to replace OSs any time soon. But AI agents are more credible as a complement that may one day become a substitute for OSs (see above, Section 4.1). In that sense, the categorisation is deserved. Two obligations – one relating to default setting and choice screens, the other relating to interoperability – should be viewed in this light.

¹⁰⁹ Own emphasis.

¹¹⁰ One could think of a scenario where a gatekeeper only allows installation of AI agents that use its on-device model. However, since foundation models do not qualify as CPS, the DMA would not prohibit such conduct.

¹¹¹ See e.g., DMA, recital 41 (distinguishing “software application, software application store or virtual assistant”).

¹¹² See e.g., Digital Markets Act, recital 49 (on using a gatekeeper “operating system, virtual assistant or web browser” to favour one’s own application).

¹¹³ *United States of America v. Microsoft Corporation*, 253 F.3d 34 (D.C. Cir. 2001).



Default Settings and Choice Screens

Article 6(3), paragraph 2, DMA holds that: “The gatekeeper shall allow and technically enable end users to easily change default settings on the operating system, virtual assistant and web browser of the gatekeeper that direct or steer end users to products or services provided by the gatekeeper.” Given that this provision is preceded by one that concerns only OSs (Article 6(3), paragraph 1, see Section 5.1) and followed by another that applies only to firms that are gatekeepers in OSs *and* in a select set of downstream services (Article 6(3), paragraph 2, part 2, see below), one might wonder about the scope of this provision. It should be interpreted as applying irrespective of whether the firm is designated in any other CPS category.¹¹⁴ Hence, should an AI agent be designated as either OS, virtual assistant *or* web browser, the gatekeeper needs to ensure that – insofar as the service defaults to other services of the gatekeeper (irrespective of whether those services have itself gatekeeper status) – it makes that **default setting changeable**. For example, if the gatekeeper provider of an AI agent would fulfil shopping requests through its (non-gatekeeper) marketplace by default, it would have to allow users to change that setting.

There is a second part to this provision goes a step further, mandating **choice screens** for a select set of CPSs. The selection is rather narrow: it only covers the *designated* “online search engine, virtual assistant or web browser to which the operating system *of the gatekeeper* directs or steers users by default, and the online search engine to which the virtual assistant and the web browser of the gatekeeper directs or steers users by default”.¹¹⁵ Hence, the obligation applies in two scenarios:

1. A firm that is *designated* as gatekeeper in both the OS CPS category *and* in the search engine, virtual assistant or web browser CPS category must show a choice screen for the designated app category at the level of the OS. In line with this obligation, Apple shows an iOS browser choice screen (it is designated in the OS and browser CPS categories) but not an iOS search engine choice screen (it is designated in the OS but not the search CPS category, as Google provides the default search engine on iOS).¹¹⁶
2. A firm designated as gatekeeper (only) in the virtual assistant or web browser CPS category needs to show a search engine choice screen. Alphabet, for example, shows a search engine choice screen on its Chrome browser (whatever the OS it is installed on).¹¹⁷

While choice screens have not always been as effective,¹¹⁸ this obligation can be meaningful in bringing AI agent choice – and thus in counteracting the foreclosure *of* AI agents by vertically integrated players (see above, Section 3.1). In particular, users may make more effective use of that choice when the market has not fully settled (i.e., there is less bias towards the familiar). Whether an AI agent choice screen is shown, however, depends entirely on its designation. If it is designated as search engine, virtual assistant or web browser, OS gatekeepers *also* designated in that category must

¹¹⁴ In line with this interpretation, Alphabet also makes its default (search) settings easily changeable on Chrome for iOS, see Alphabet, ‘EU Digital Markets Act Compliance Report: Non-Confidential Summary’ (7 March 2025) <https://transparencyreport.google.com/report-downloads>.

¹¹⁵ Digital Markets Act, art 6(3), paragraph 2, part 2 (own emphasis).

¹¹⁶ Apple’s Non-Confidential Summary of DMA Compliance Report, 7 March 2024, <https://www.apple.com/legal/dma/dma-ncs.pdf>. For its browser choice screen, see <https://developer.apple.com/support/browser-choice-screen/>.

¹¹⁷ Alphabet, ‘EU Digital Markets Act Compliance Report: Non-Confidential Summary’ (7 March 2025) <https://transparencyreport.google.com/report-downloads>.

¹¹⁸ For an overview, see Omar Vasquez Duque, ‘Taking Behavioral Antitrust Seriously: On Default Agreements as Exclusive Dealing and the Debiasing Potential of Default Randomization’ (2024) 84 Maryland Law Review 143, 191–197.



show a choice screen. This means that Microsoft (Windows), Apple (iOS) and Google (Android) would have to provide search/assistant/browser choice screens on their OS if their AI agent is designated in one of those three categories.¹¹⁹ If a firm's AI agent is designated as virtual assistant or web browser, it needs to show a search engine choice screen (but not a choice screen for other virtual assistants or browsers).

There are three problems with this obligation in an AI agent context. First, a potential doubling up could potentially nullify certain obligations. If a firm's AI agent is designated as OS – a distant but not unthinkable possibility (see above, Section 4.1) – it might need to show a search/assistant/browser choice screen. However, an assistant choice screen is obviously without purpose in this scenario. Further, if the firm's (traditional) OS is also designated – i.e., it has two OS designations – it would not have to show any choice screen. This makes technical sense – if the agent essentially overlaps with the OS, you cannot just swap it out – but it may limit contestability. By contrast, a search or browser choice screen may be interesting here, especially as agents make great use of these adjacent services. Again, however, it may be technically difficult to swap out the search index the agent relies on.

Second, the more meaningful obligation that allows for an AI agent (not just a search engine) choice screen, is limited to vertically integrated firms that are gatekeepers at two levels: upstream (OS) and downstream (AI agent, whether designated as search engine, virtual assistant or web browser). This means that it does not apply when an OS gatekeeper contracts out its provision of an AI agent. Apple, for example, has partnered with OpenAI to provide its agent (with another partnership with Google in the works).¹²⁰ Insofar as those agents are not considered its own, it would escape any choice screen obligation on iOS even if the agents rise to the level of gatekeeper status.

Third, AI agents may communicate with each other, with one agent invoking another, possibly more specialised, AI agent to perform a task (see above, Section 2). However, if an AI agent is designated as a virtual assistant (or web browser), it does itself not have any obligation to offer a choice screen for which other AI agents it may invoke per default. Generally, this issue seems of lesser concern, however, as it may become very complex quickly if a user were always required to mix-and-match the ecosystem of AI agents she wants to use.

Interoperability

Article 6(7) DMA centres on interoperability and has two parts to it.

1. The first part of the provision obliges firms with a designated OS or virtual assistant to allow other hardware and service providers free and effective interoperability with hardware and software features as are available to services or hardware provided by the gatekeeper.

¹¹⁹ For the OS designations, see *Microsoft – online social networking services, number-independent interpersonal communications services and operating systems* (Cases DMA.100017, DMA.100023, DMA.100026) Decision of 5 September 2023, *Apple – online intermediation services (app stores), operating systems and web browsers* (Cases DMA.100013, 100025, 100027) Commission Decision of 5 September 2023 and *Alphabet – OIS verticals, OIS app stores, online search engines, video sharing, number-independent interpersonal communications services, operating systems, web browsers, online advertising services* (Cases DMA.100011, 100002, 100004–100006, 100009, 100008 and 100010) Decision of 5 September 2023.

¹²⁰ 'OpenAI and Apple announce partnership to integrate ChatGPT into Apple experiences' (*OpenAI*, 10 June 2024) <https://openai.com/index/openai-and-apple-announce-partnership/>; Mark Gurman, 'Apple Is in Talks to Let Google Gemini Power iPhone AI Features' (*Bloomberg*, 18 March 2024) <https://www.bloomberg.com/news/articles/2024-03-18/apple-in-talks-to-license-google-gemini-for-iphone-ios-18-generative-ai-tools>.



2. The second part of the provision obligates firms with a designated OS (so not virtual assistant) to allow business users and other providers of services (not hardware) that integrate with CPSS free and effective interoperability with OS, hardware and software features as are available to the gatekeeper when providing such services.

The goal of interoperability is that “all elements of hardware or software work with other hardware and software and with users in all the ways in which they are intended to function”.¹²¹ The degree of interoperability to be provided is determined by a non-discrimination standard: other providers must receive the *same* interoperability affordances as the gatekeeper’s own hardware and services – hence, the obligation is one of **equal vertical interoperability**. The provision is set up to prevent situations of input foreclosure, where an OS/virtual assistant gatekeeper uses interoperability to discriminate in favour of (self-preference) its own downstream services to the detriment of competitors (see above, Section 3.1). There thus needs to be a degree of vertical integration (OS and downstream service) for the obligation to apply.

In the context of AI agents, the obligation can play a dual role. First, it is crucial in ensuring that the providers of third-party AI agents can compete on a level playing field with the OS provider’s first-party AI agent (part 2). On Android, for example, the AI agents of third parties should get the same access as Google’s own agent. There is a potential problem with building this obligation around vertical integration: if the OS provider contracts out the provision of an AI agent, are we then still talking about a “service provided by the gatekeeper”? Insofar as the OS gatekeeper’s AI agent is provided *in partnership with* a third party, such as in the case of Apple and OpenAI (for Apple Intelligence), the obligation would still seem to apply. But if the OS gatekeeper pre-installs an AI agent provided entirely by a third party (perhaps in exchange for a revenue share), the equal interoperability obligation would seem not to apply.

Second, AI agents themselves can also be covered by the obligation if they are designated as either OS or virtual assistant (part 1). In that case, their providers do not benefit from interoperability but have to grant it to others. If the AI agent provider manufactures a connected device to run, support or integrate with the agent, for example, it would have to grant competing hardware providers equal interoperability. As seen in a specification decision, such interoperability could relate to the display of notifications, proximity-triggered pairing, file transfer, wi-fi connection, near-field communication control, and so on.¹²² It would also require the gatekeeper to institute a process to handle interoperability requests.¹²³

5.3. Obligations for Search Engines

When it comes to search engines, the obligations relate to self-preferencing, data sharing, and fair, reasonable, and non-discriminatory (FRAND) terms. As we explain below, however, that the self-preferencing obligations may extend beyond search engines.

¹²¹ Digital Markets Act, art 2(29).

¹²² *Apple – Operating Systems – iOS – Article 6(7) – SP – Features for Connected Physical Devices* (Case DMA.100203) Commission Decision of 19 March 2025.

¹²³ For an example, see *Apple – Operating Systems – iOS – Article 6(7) – SP – Process* (Case DMA.100204) Commission Decision of 19 March 2025.



Self-preferencing in Ranking

Article 6(5) DMA prohibits self-preferencing: “The gatekeeper shall not treat more favourably, in ranking and related indexing and crawling, services and products offered by the gatekeeper itself than similar services or products of a third party.” The wording of the prohibition – in particular, its reference to ranking, indexing and crawling – would seem to restrict it to search engines, which is probably due to the prohibition’s origins in the *Google Shopping* case (see above, Section 3). If an AI agent is designated as a search engine, this is not a problem. If an AI agent is designated in another CPS category, however, the provision’s terminology would need some stretching in order to apply, though this does need not make the prohibition inapplicable.

For instance, an AI agent does not necessarily “rank” services and products in the sense that it provides a ranking to the user. Even though an agent may rank options for a user, it could also directly buy, subscribe or more generally consume on the user’s behalf. The DMA, however, subscribes to a wide definition of “ranking”:

The relative prominence given to goods or services offered through online intermediation services ... or virtual assistants, or the relevance given to search results by online search engines, as presented, organised or communicated by the undertakings providing online intermediation services ... virtual assistants or online search engines, irrespective of the technological means used for such presentation, organisation or communication and irrespective of whether only one result is presented or communicated.¹²⁴

This wide definition may have been adopted in view of Google’s gradual evolution from a *search* engine (with different results – the proverbial “ten blue links”) to an *answer* engine (with just one [prominent] result), or perhaps to cover virtual assistants, which tend to skip the ranking of different options and are explicitly included here. Given this definitional exercise, Article 6(5)’s self-preferencing prohibition would therefore seem to cover AI agents when designated as search engines or, indeed, in other CPS categories such as virtual assistants.

The second part of the provision, which obliges the gatekeeper to “apply transparent, fair and non-discriminatory conditions to such ranking” is also difficult to apply in the context of AI agents. Even if there is ranking going on (behind the scenes), it is a very challenging task to provide “transparency” about how an AI agent (typically based on a deep neural network) derives decisions and suggestions from its training and the available data. While explainable AI (XAI) is an active area of research, explanations need to be on a high level of abstraction from the underlying neural network, and it is therefore doubtful whether they are really suitable to determine with sufficient legal certainty whether “fair and non-discrimination” conditions have been applied. The problem here is thus one of implementation and enforceability.

Data Sharing

A second obligation applicable to search engines is found in **Article 6(11) DMA**: “The gatekeeper shall provide to any third-party undertaking providing online search engines, at its request, with access on fair, reasonable and non-discriminatory terms to ranking, query, click and view data in relation to free and paid search generated by end users on its online search engines.” More clearly than for Article

¹²⁴ DMA, art 2(22) (eliding the references to social networks and video-sharing services).



6(5) DMA, the addressees of this **search data sharing obligation** are limited to gatekeeper search engine providers. In moving from traditional search engines to AI agents, however, the scope of the covered data loses some clarity. “Query” data is still straightforward enough, as AI agents require input, even if it may come in a wider variety of forms (text, voice, visual...). With regard to “ranking”, we pointed out the issue above. “Click and view” data, finally, refers to the typical follow-up actions after a search engines shows its ten blue links. Interactions with an AI agents may but certainly not need to result in a click – the whole point of moving from search to answer engines is to obviate the need for such a click. Similarly, the user of an AI agent may not “view” any results at all – the action may simply be carried out (or the answer might be “heard” rather than viewed). This interpretation difficulty can be resolved by focussing on click and view data at one level of abstraction – so on follow-up actions generally rather than on specific methods of navigating the CPS – but even then, there is still the practical difficult of sharing certain data (e.g. ranking data).

Finally, the question is how useful the shared data would be for contestability in the market for AI agents. Like in the case of search engines, also for AI agents the answer to this question likely depends to a large extent on the details, such as the granularity and extent of the data provided, and the degree of anonymisation that is done. Generally, as we have discussed in Section 2, AI agents can improve and learn from user feedback, and thus, access to query, click and view data seems to also be, in principle, meaningful in the context of agentic AI. However, given the broader scope and more complex operations that agentic AI can perform, it seems challenging (albeit not infeasible) to devise datasets to be shared with competitors that are both meaningful and preserve anonymity, as demanded by the DMA. This has already proven to be a challenging task in the enforcement of data sharing of search engines, for which the provision was designed, and where both the scope and the user-data generated are well defined.¹²⁵ Yet, the datasets to be provided under Article 6(11) DMA have been criticised by competitors for leaving out most of the (relevant) data.¹²⁶

FRAND Terms

A final obligation for search engines (as well as app stores and social networks) is found in **Article 6(12) DMA**, which obliges them to apply “fair, reasonable, and non-discriminatory general conditions of access for business users”. Such **FRAND obligations** are well-known from standard-setting procedures, but they remain a controversial and hard-to-define concept.¹²⁷ The obligation of “non-discriminatory” conditions is similar to the ban on self-preferencing and thus comes with similar difficulties (see above). What is “fair and reasonable” needs to be established in relation to the service in question, i.e., AI agents. The issue of enforceability returns here, as the (background) operation of an AI agent is more shielded from the view of users and authorities than that of a traditional search engine.

The implementation of the above, search-engine-specific obligations were already not always as easy to apply to traditional search. Google’s implementation of the self-preferencing prohibition of Article 6(5), which is also in 6(12), already led to a non-compliance investigation – even after experience with a similar non-discrimination remedy under competition law.¹²⁸ Gatekeepers in scope of Article 6(12)

¹²⁵ For a discussion of the complexity of sharing search data, see, e.g., Krämer and Schnurr (n 22), Section III.B.4; Inge Graef and Jens Prüfer, ‘Governance of data sharing: A law & economics proposal’ (2021) 50 Research Policy.

¹²⁶ See e.g., ‘Roadblocks to Competition: Investigate Google’s Non-Compliance with the EU’s Digital Markets Act’ (*DuckDuckGo*, 20 November 2024) <https://spreadprivacy.com/investigate-google-dma/>.

¹²⁷ See e.g., Yann Ménière, ‘Fair, Reasonable and Non-Discriminatory (FRAND) Licensing Terms: Research Analysis of a Controversial Concept’ (JRC Science and Policy Report, 2015).

¹²⁸ *Alphabet – Online Search Engine – Google Search – Article 6(5)* (Case DMA. 100193) Commission Decision of 25 March 2024.



generally reported they already complied with its FRAND obligation, which the European Commission has not yet questioned.¹²⁹ As described above, also the sharing of search data has proven to be challenging. The implementation and enforcement of these obligations only becomes more complex for AI agents.

5.4. Obligations Independent of CPS Category

Finally, the DMA contains a number of obligations that apply to all CPSs, independent of their category. These obligations can be grouped, roughly, in those putting restrictions on data use, those mandating data portability, and those relating to contractual obligations and restrictions.

Restrictions on Data Combination and Use

A series of obligations under **Article 5(2) DMA** are meant to discourage harvesting and cross-use of personal data, unless the user has explicitly consented. These provisions, originating from the *Bundeskartellamt v. Facebook* case, are in fact, rooted in the General Data Protection Regulation, and could arguably also be enforced by the responsible data protection authority. The DMA has included these provisions because many of the gatekeepers employ ecosystems of services, which presents them with a competitive advantage over smaller rivals, not having comparable ecosystems, to collect, combine and cross-use personal data across services.

Generally, restricting the collection and use of data will also limit the economic efficiencies that are associated with learning from data. Everything else being equal, enabling data sharing (which shares those efficiencies with rivals) is therefore to be preferred over putting restrictions on data use.¹³⁰ However, given the complexities of data sharing discussed above, data use restrictions may be second-best. In this vein, these obligations also seem to make sense in the context of AI agents, which are also likely being deployed within ecosystems of services of the same provider, including, e.g., OSs, browsers, search engines, marketplaces and so on.

Further, **Article 6(2) DMA** restricts the use of “non-public data in competition with business users that was gathered through the activity of business user during their use of the relevant core platform services.” Albeit designed with marketplaces in mind, this obligation may also be reasonable in the context of AI agents. Specifically, as the AI agent interact with many third-party software applications, which make themselves available for such interaction through offering tools and APIs (e.g., like in the example of Apple’s App Intents¹³¹), the agent can potentially learn about user preferences and user behaviour – as the user interacts through the agent with the third-party applications – that the provider of the agent may use in competition with those third-party app providers. In similar spirit, with the integration of AI into OSs, OS providers are already collecting data about the use of each individual app.¹³²

¹²⁹ See e.g., Alphabet, ‘EU Digital Markets Act Compliance Report: Non-Confidential Summary’ (7 March 2025) <https://transparencyreport.google.com/report-downloads>.

¹³⁰ Compare Jan Krämer and Shiva Shekhar, ‘Regulating Digital Platform Ecosystems Through Data Sharing And Data Siloing: Consequences For Innovation And Welfare’ (2025) 49 MIS Quarterly 123; Krämer and Schnurr (n 22).

¹³¹ <https://developer.apple.com/documentation/appintents/app-intents>.

¹³² <https://talk.tidbits.com/t/ios-18-what-does-learn-from-this-app-mean-with-siri/30081>.



Data Portability

Independent of the CPS type, designated CPS must provide “end users and third parties authorised by an end user” as well as business users effective portability of data according to **Articles 6(9) and 6(10) DMA**. Specifically, Article 6(9) relates to “data provided by the end user or generated through the activity of the end user in the context of the use of the relevant core platform service” and access to such data has to be provided continuously and in real-time. As we established above (see Section 2), access to such data provided by end-users may be very useful for the training and optimisation of AI agents (and assistants). A popular AI agent has the ability to collect more conversational data and generally more interactions with the user, which can be used to improve the AI agents’ outcomes, through reinforcement learning and other means. In addition, as AI agents are likely to have persistent memory and greater ability to learn and adapt to the user, data portability can counteract data-induced switching costs. Thus, the provision in Article 6(9) is generally considered to be very useful to promote market contestability also in the context of agentic AI.

By contrast, the provision in Article 6(10) does not seem to fit the context of AI agents well. The provision demands that “the gatekeeper shall provide business users ... aggregated and non-aggregated data, including personal data, that is provided for or generated in the context of the use of the relevant core platform services”. This provision was devised in the context of online marketplaces, where business users’ interaction with end users is fully mediated by the platform. The difference to the AI agent context is that the agent acts as a surrogate for the user. Agents thus have the potential to disintermediate the platform and act with business users directly. In this case, there would be no information asymmetry introduced by the platform that is to be corrected through data portability. In the case where AI agents interact with business users via the platform or marketplace, then it is again only the platform (and not the agent) that can meaningfully provide such data portability.

In any case, even when limiting our attention to Article 6(9), the details of the implementation of the provision can become quite complex in the context of AI agents. Even in the seemingly simpler contexts of “traditional” CPSs, such as social networking, search or messaging, gatekeepers have interpreted “real-time and continuous” data portability as allowing users to schedule daily downloads (once every 24 hours) for a limited period of time.¹³³ Furthermore, there still exists some legal uncertainty as to the scope of data portability, i.e., what data precisely was “provided by ... and generated through the activity of the end user”. In order to be meaningful for purposes of contestability, a wide interpretation of the scope of portable data would be helpful.¹³⁴ In the context of agentic AI, it is evident that, at a minimum, queries to the AI (prompts) are covered by the data portability provision, but reasonably also the responses by the AI are legitimate “data generated in the context of use” that become subject to portability. In this way, users would be able to port whole conversations with the AI. On the one hand, this would facilitate contestability, as conversational datasets would become portable.¹³⁵ On the other hand, widening the scope of data portability may undermine innovation incentives as it promotes the possibility of “distillation”, i.e., a process where

¹³³ See, e.g., the compliance reports of Alphabet, Apple and Meta (available at <https://digital-markets-act-cases.ec.europa.eu/reports/compliance-reports>), all of which apply similar restrictions.

¹³⁴ See Jan Krämer, ‘Personal Data Portability In The Platform Economy: Economic Implications And Policy Recommendations’ (2021) 17 *Journal of Competition Law & Economics* 263.

¹³⁵ First efforts to support such data portability of conversational data sets are already underway, see Chris Riley, ‘The future of AI hinges on data portability and APIs’ (*Data Transfer Initiative*, 11 February 2025) <https://dtinit.org/blog/2025/02/11/future-of-AI-portability>.



AI assistants/agents are being trained by extracting data from other models, allowing them to catch up to the incumbent, but without undergoing the same efforts.¹³⁶ Nevertheless, along with other observers¹³⁷ we suggest that those responses should be included in the portable dataset. Likewise, data observed by the AI, like sensor data or location, should be subject to data portability. This would also be in line with the provisions in the Data Act concerning connected products.¹³⁸

However, there are also inherent limits to data portability. AI models are trained and possibly adapted to users' preferences through finetuning of billions of parameters. From today's vantage point, it does not seem feasible nor sensible to port parameter values (or a selection thereof) from one provider to another. Given the amounts of investment that go into training a model, this approach would also not be proportionate.¹³⁹

Notably, the data portability provision under the DMA can also apply to firms designated as "emerging gatekeepers".

Contractual Obligations and Restrictions

A final set of DMA obligations that apply independent of CPS category relate broadly to contractual obligations and restrictions. As we explain below, however, they have sometimes been adopted with a specific scenario in mind and are not always easily translated beyond that scenario.

In accordance with **Article 6(13) DMA**, a gatekeeper must maintain proportionate termination conditions, ensuring that termination can be exercised without undue difficulty. **Article 5(6) DMA** prevents gatekeepers from having business users raise compliance issues with authorities and courts. **Article 5(3) DMA** bans so-called "most-favoured-nation clauses", which have been maintained by online intermediation services (e.g., hotel booking or marketplace platforms). Insofar as an AI agent intermediates (e.g., for its end-user's purchases), its provider cannot (contractually) prevent business users from offering the same products/services at better conditions elsewhere.

Some obligations relate to payments and associated access, and are therefore also mainly relevant with regard to a platform's intermediation function. **Article 5(4) DMA** prohibits gatekeepers from maintaining anti-steering policies. Hence, a designated AI agent provider engaging in intermediation could not prevent business users from directly communicating and transacting with end-users. **Article 5(5) DMA** backs up that provision by obliging gatekeepers to let their end-users access and use, through its CPSs, services/content/subscriptions/features acquired outside of the gatekeeper's CPS. In an AI agent context, this provision gives end-users freedom to invoke related services of their choice, e.g., to specifically buy via the marketplace where they have an account/subscription, rather than the one suggested by the AI agent.

In a similar vein, **Article 5(7) DMA** prohibits gatekeepers from pushing certain services on their business and end-users. The covered services are not CPSs themselves but often form part of a gatekeeper's larger ecosystem. In further support of Articles 5(4)–(5), the provision ensures that end-

¹³⁶ See <https://www.theverge.com/news/601195/openai-evidence-deepseek-distillation-ai-data>.

¹³⁷ See Chris Riley, 'Digging in on personal AI portability' (*Data Transfer Initiative*, 4 June 2024) <https://dtinit.org/blog/2024/06/04/digging-in-personal-AI>.

¹³⁸ Regulation (EU) 2023/2854 of the European Parliament and of the Council on harmonised rules on fair access to and use of data, OJ L2023/2854, arts 3–5.

¹³⁹ For a similar view, see <https://dtinit.org/blog/2024/06/04/digging-in-personal-AI>.



users can *use*, and business users can *interoperate with*, “a payment service, or technical services that support the provision of payment services, such as payment systems for in-app purchases” of choice. The same goes for identification systems. Freedom must also be provided with regard to browser engines. This obligation was adopted for OS gatekeepers that mandate the use of their own browser engine, also for competing browsers. In accordance with the provision, an OS gatekeeper with a browser engine could not prescribe its use by the providers of competing AI agents either. The provision becomes difficult to apply when moving beyond the OS and online intermediation (e.g. app stores, marketplaces) CPS category, even if it theoretically applies there. Consider, for example, a single designation (of an AI agent or different service) as web browser: would its provider also be obliged to offer its users a choice of engine in that case? Notably, in our context, the provision in Article 5(7) does not (yet) guarantee freedom of choice regarding AI agents (or virtual assistants). Instead, it is explicitly restricted to identification services, payment services, and web browser engines. We will elaborate on this point below (Section 6).

Article 5(8) DMA prohibits tying one gatekeeper service (or even a non-designated service that meets the user thresholds; for those, see above, Section 4.2) to another gatekeeper service. Hence, a gatekeeper cannot use access to one service (“tying” service) to reinforce another (“tied” service). The scope of the provision is limited, however, by the fact that the tied service must already have gatekeeper status (or meet the user thresholds); it does not apply to a nascent service. Nevertheless, it can prove useful for defensive leveraging scenarios, where a gatekeeper would want to maintain its strong position in a CPS category by tying that service to one that is in an *even* stronger position. Depending on its strength, an AI agent could serve as tying or tied product. In the former case, the concern would be one of foreclosure *by* an AI agent; in the latter, *of* an AI agent (see above, Section 3).

Article 6(6) DMA, finally, prohibits gatekeepers from restricting the ability from end-users to switch between different apps accessed via their CPSs. The scenario targeted here is one where a gatekeeper in control of an OS or middleware CPS makes user switching between apps more difficult, presumably in favour of the (gatekeeper’s) default service that is pre-installed. Hence, the provision could promote switching from the gatekeeper’s to a third-party AI agent – in other words, it prevents foreclosure *of* an AI agent (see above, Section 3.1). If the AI agent itself is designated, however, the provision would imply that the gatekeeper cannot restrict switching between further downstream services (e.g. marketplaces) – in other words, it prevents foreclosure *by* an AI agent (see above, Section 3.2).



6. Should the DMA be Revised in the Face of Agentic AI?

In this section, we tackle the difficult question of whether the DMA should be revised in the face of agentic AI. Answering that question is an exercise in synthesis: we bring together the competitive concerns (Section 3), checking to which extent they are addressed by obligations in the DMA (Section 5) as they might apply to AI agents, which often depends on their CPS qualification (Section 4). The structure follows the two sets of competitive concerns identified earlier: foreclosure of an AI agent (Section 6.1) vs foreclosure by an AI agent (Section 6.2), with the former being a more immediate and the latter being a more distant concern (time-wise).

6.1. Foreclosure of an AI Agent

In the context of foreclosure of an AI agent, we will focus on the most likely scenario, where the gatekeeper of an existing OS, designated under the DMA, pre-installs or otherwise integrates an AI agent on the OS, but seeks to foreclose the use of a third-party AI agent on that OS. Arguably without a pre-installed AI agent, the OS provider would have no incentive to foreclose a third-party AI agent. The pre-installed AI agent may or may not be designated under the DMA in its own right. If the AI agent is designated, we assume it is designated as a “virtual assistant”, which seems to be the most fitting category under the DMA. In addition, we assume that the pre-installed AI agent is considered the gatekeepers “own” service, which in our understanding would – or at least should – legally also be considered the case if the agent is provided through a close industry cooperation. A separate issue, not addressed here, are contracts between a gatekeeper and another player that involve potentially large payments for pre-installing or defaulting to a third-party AI agent.¹⁴⁰

As noted in Section 3, pre-installation of an “own” AI agent on the OS can give competitive advantages to the pre-installed agent over third-party agents: users may be reluctant to switch to another AI agent if it is not possible to uninstall the pre-installed one, or if it is not possible to set the new AI agent as the default. Depending on the circumstances, pre-installation can be viewed as a form of self-preferencing (one’s own software over others) or a form of tying (shipping one’s own software with the OS). When the OS on which the AI agent runs is designated, the DMA’s obligations in regard to uninstallation and default settings (Article 6(3)) as well as switching (Article 6(6)) and tying (Article 5(8)) generally apply, but the extent to which depends on whether pre-installation concerns the gatekeepers own AI agent, and whether this agent is designated as a virtual assistant CPS in its own right.

Another form of self-preferencing may come in the form of giving the pre-installed own AI agent more comprehensive access to the OS, i.e., to the hardware or software features of the device, including more capabilities to collect data on the device (e.g., about the users’ location or usage of other apps). Subject to strictly necessary and proportionate integrity measures, Article 6(7) DMA comprehensively addresses these concerns, i.e., the gatekeeper has to provide equal vertical interoperability to other AI agents, as they provide for their own, and irrespective of whether the own agent is designated as a

¹⁴⁰ In the spirit of, for example, the case where Google is paying Samsung an “enormous sum” to pre-install Gemini, see <https://www.theverge.com/news/652746/google-samsung-gemini-default-placement-antitrust-trial>.



CPS in its own right. Indeed, Article 6(7) would apply even irrespective of whether the pre-installed AI agent is deemed an essential part of the OS or not, as the obligations only refers to “features” available to the gatekeeper. However, provided that AI agents require far-reaching and deep access to the OS to fulfil their tasks in a meaningful way, allowing equal interoperability is likely to become a contentious issue in enforcement. The closest comparable scenario in today’s enforcement actions of the DMA is probably that of enabling alternative browser engines on iOS. However, enabling alternative AI agents is likely to be far more complex, also involving even stronger trade-offs for security and integrity due to their ability to react to user input in various forms (e.g., through visual, audio or motion cues), and their ability perform tasks autonomously on behalf of the user. Therefore, the API access necessitated by AI agents would not only involve those needed by an alternative browser engine (which is required by the AI agents in order to browse the web autonomously), but also additional ones.

The DMA also includes comprehensive obligations to equate the ability of first-party and third-party service providers to harness data and to learn from data. This includes the provisions that aim at reducing ecosystemic benefits from data combination and cross-use (Articles 5(2) and 6(2)) as well as the prevention of data-induced lock-in effects through data-portability (Article 6(9)–(10)). These all apply in the context of AI agents shipped with an OS; however, the data portability provisions of the DMA only apply to the pre-installed AI agent directly if it is designated as CPS itself (unlike the data portability provision under Article 20 GDPR, which applies in any case).

Notably, the data sharing provisions in relation to “click and query” data under Article 6(11) DMA would likely not apply in the present context, unless the pre-installed AI agents is designated as a search engine. Thus, representative, anonymised data (across users) on queries and prompts, as well as user’s reactions (“clicks”) would not be available to alternative AI agents, even if the pre-installed AI agent is deemed a virtual assistant CPS, and even though we think this provision would generally be useful in this context (see above, Section 5.3).

Another form of anti-competitive conduct could revolve around leveraging market power from the OS (or access to the CPS app store for the OS¹⁴¹) to promote the use of the pre-installed AI agent in relation to business users. For example, the OS provider could require that software developers (i.e. business users of the OS) integrate features in their software so they work seamlessly with the AI agent (e.g., APIs that facilitate execution of tasks in the app by the agent). This would provide an advantage to the AI agent over other agents that do not benefit from such forced integration. Especially providers of popular apps that are themselves seeking to offer AI agents may be reluctant to integrating competing AI agents in their products.¹⁴² Article 5(7) DMA currently does not cover this case, because it is limited to forced integration with identification services, payment services and web browser engines of the gatekeeper. A potential improvement of the DMA in this regard could therefore be to extend Article 5(7) to forced integration with “virtual assistants”, whether yet designated as CPS or not. In case the pre-installed AI agent is already designated as CPS, Article 5(8) DMA could also apply here to some as extent, as far as business users (or end-users) are forced to “subscribe to, or register with” the pre-installed AI agent. However, it is questionable whether this can

¹⁴¹ In this context, the DMA’s provisions on the right of “installation” (see above, Section 5.1) become relevant, as they limit the extent to which such leverage of access to the OS onto other services of the gatekeeper can be exercised.

¹⁴² Consider, for example, the case of Meta which does not want to integrate features of Apple Intelligence into their products, see <https://9to5mac.com/2025/04/16/meta-blocks-apple-intelligence-on-facebook-and-its-other-ios-apps/>.



be extended to forced “integration with” the gatekeeper’s AI agent. Further, it is also questionable whether the FRAND access provision in Article 6(12) DMA provide any reasonable protection to business users not willing to integrate with the gatekeeper’s AI agent, as this applies only to “application stores, online search engines and online social networking services”, but not to OSs.

A final anti-competitive conduct to be considered here is that of self-preferencing in ranking. While Article 6(5) DMA was designed in reference to search engines, we argued above (see Section 5.3) that the obligation is not strictly limited to this context and may well apply also to the virtual assistant CPS category. This may matter in the present context insofar an OS provider, who ships the OS with a pre-installed AI agent but is also operating a CPS search engine may want to demote rival AI agents in the search results of that search engine. Similar arguments can be made for the findability of alternative AI agents in app stores or, indeed, in search functions implemented on the OS itself.

Table 1 below summarises the applicability of the DMA provision in the context of foreclosure of an AI agent.

A final set of anti-competitive concerns mentioned in Section 3.1 but not addressed here revolve around foreclosure of off-device bottleneck inputs for the development and functioning of AI agents, such as foundational model, training data and (cloud) computing resources. Indeed, the DMA has little to say about this. AI models cannot be reasonably mapped to any of the existing CPS categories, nor is there an obligation to share training data of popular services by gatekeepers more generally. Data sharing is limited to click and query data by search engines and to data portability by users. Even though cloud computing services are a CPS category, and Article 6(12) DMA obliges gatekeepers to apply FRAND access terms for business users, this provision is strictly limited to search engines, app stores and social networking services – it does not apply to cloud computing services. Thus, as it stands, the DMA would not prevent a cloud computing provider, if a firm would be designated as such in the future, from discriminating against specific AI agent providers (albeit competition law would, of course, still apply in this context).

Table 1: DMA obligations that apply to a designated OS provider with pre-installed AI agent, mapped to potential competition concerns relating to the foreclosure of a third-party AI agent

DMA OBLIGATION FOR DESIGNATED OS WITH PRE-INSTALLED AI AGENT	INPUT FORECLOSURE OF AI AGENT THROUGH DENIAL OF EQUAL ACCESS TO ...		DISTRIBUTION FORECLOSURE OF AI AGENT THROUGH ...		
	data	on-device software & hardware features	pre-installation	self-preferencing	tying
uninstallation (6(3), para 1)			✓	(✓)*	(✓)*
changeable default/choice screen (6(3), para 2)			✓ ¹	(✓)*	(✓)*
equal vertical interoperability (6(7), part 2)	✓ ²	✓	✓	(✓)*	(✓)*
end-user switching (6(6))			✓	(✓)*	(✓)*
no tying (5(8))					✓ ³



data portability (6(9)–(10))	✓ ³				
restrictions on data use (5(2) & 6(2))	✓				
no self-preferencing in ranking (6(5))				(✓) ⁴	
sharing of click and query data (6(11))	✗ ⁵			✗ ⁵	
no forced integration with gatekeeper services (5(7))	✗		✗	✗	✗
FRAND access for business users (6(12))				?	

Notes:

- * Insofar as pre-installation is considered a form of self-preferencing and/or tying.
- 1. Obligation only applies if gatekeeper *also* provides *own* virtual assistant (not for third-party pre-installed assistants): if gatekeeper's virtual assistant also designated, choice screen must be offered; otherwise, default settings must only be changeable.
- 2. In relation to on-device data (e.g. GPS data; usage data of installed apps, as it is available to the OS and made available to gatekeeper's own virtual assistant).
- 3. Only if the pre-installed virtual assistant is also designated as CPS.
- 4. In relation to demotion of the third-party AI agent in search results provided by the OS or, if designated, pre-installed AI agent; and – if applicable – in the gatekeeper's designated search engine and online application store.
- 5. Unless the AI agent is designated as "search engine".

6.2. Foreclosure by an AI Agent

Foreclosure of other services *by* an AI agent is a more distant concern than foreclosure *of* an AI agent. This concern only manifests when a provider of an AI agent – whether one of the current gatekeepers or a new player – becomes dominant (or, in the words of the DMA, an important gateway) in its own right. At that point, it may be able to (ab)use that position to foreclose the providers of services downstream of the AI agent. It must have an incentive to do so, which may come from the fact that it provides some of those downstream services itself or that it has agreements granting preferential status to the services of third parties. This concern, of a "super-gatekeeper" that a variety of third-party services rely on was already suggested for voice assistants but seems more realistic in a world of capable AI agents intermediating a wide range of tasks for their users.

Here, more so than in the previous section, the CPS category in which the AI agent is designated with gatekeeper status matters. Recall that four categories provide a potential fit: virtual assistant, web browser, and search engine in the short- to medium-term; OS in the long-term (see above, Section 4.1). Let us go over the relevant DMA obligations and the CPS categories they apply to, as represented in the table below.



Table 2: DMA obligations that address foreclosure by an AI agent, mapped to potential CPS categories in which the AI agent may be designated

OBLIGATION	OS	ASSISTANT	BROWSER	SEARCH
no self-preferencing in ranking (6(5))	(✓)	✓	(✓)	✓
FRAND access for business users (6(12))	✗	✗	✗	✓
changeable default (6(3), para 2)	✓	✓	✓	✗
no tying (5(8))	✓	✓	✓	✓
search choice screen (6(3), para 2)	✗	✓	✓	✗
equal vertical interoperability (6(7), part 1)	✓	(✓)	✗	✗
free app installation (6(4))	✓*	✗	✗	✗
no use of business user data (6(2))	✓	✓	✓	✓
free communication & no forced integration with gatekeeper services (5(4), 5(7))	✓	✓	✓	✓

* the AI agent of the OS gatekeeper must (also) be designated in its own right

A large number of the obligations can be grouped under the “self-preferencing” heading. Leading this group is Article 6(5) DMA’s prohibition of favourable treatment in “ranking and related indexing and crawling”. While its terminology clearly originates from search engines, the DMA’s broad definition means that virtual assistants (and, if need be, even OSs and web browsers) are covered. The provision also mandates applying “transparent, fair and non-discriminatory conditions” to such ranking. However, providing transparency is technically tricky with regard to AI agents, which also means that establishing fairness and non-discrimination become difficult. Moving on, the FRAND obligation of Article 6(12) DMA applies only to AI agents designated as search engines. There is a partial overlap here with Article 6(5), which covers non-discrimination in ranking. Unless an AI agent is designated as search engine, however, it need not also ensure its access terms are fair and reasonable under Article 6(12).

Another form of self-preferencing is covered by Article 6(3), paragraph 2, DMA: If the provider of a designated AI agent uses it to steer users to its other services (e.g., its marketplace) by default, then it must make that default setting changeable. This obligation only applies, however, if the agent is designated as OS, virtual assistant, or web browser – it does not apply to search engines. The obligation can be seen as complementary to Article 6(5), which does cover search engines, preventing them from steering users to specialised search services by default. While the self-preferencing that motivated the respective obligations – app defaults vs web results – is different, the obligations are formulated so as to cover a wider range of favourable treatment. The important difference is in the exact remedy: While self-preferencing in search is excluded, app defaults only require changeable settings. The question is how many users will take the step to do change them. There is also a relation with Article 5(8) DMA, which prohibits tying between gatekeeper services. An AI agent could function as tying service, to prop up a different service with gatekeeper status (the tied service), in which case



we are talking about foreclosure *by* an AI agent, but the requirement of two gatekeeper services limits the scope of this provision, even if it does not matter in which category they are designated.

Cognisant of users' *status quo* bias, lawmakers went a step further than changeable defaults by mandating choice screens for certain (constellations of) CPSs in the second part of Article 6(3), paragraph 2, DMA. While choice screens can remedy foreclosure of an AI agent to some extent (see above, Section 6.1), they have little impact on foreclosure by an AI agent. The obligation only kicks in when an AI agent is designated as virtual assistant or web browser, in which case it would need to show a choice screen but only for search engines. As we discussed above (Section 4.1), AI agents do tend to integrate (with) search engines, so the obligation may have merit, even if it is difficult to implement due to the depth on integration with a search engine (and in particular the underlying index).

A final form of self-preferencing, via interoperability, is targeted by Article 6(7) DMA, and may also be relevant for foreclosure by an AI agent. If an AI agent is designated as either OS or virtual assistant, it needs to provide equal vertical interoperability to other hardware and software providers. A possible scenario here may be that in which the gatekeeper provider of an AI agent also provides a connected device for/with the agent, in which case it would have to grant competing hardware providers the same level of interoperability. As AI agents become system(-like) software, interoperability with the agent's software features likely also becomes increasingly important. Note, however, the second part of Article 6(7), which is most important for such access, applies only to OS gatekeepers.

A number of other obligations are more *sui generis*¹⁴³ and have a (theoretically) broad scope. First, should an AI agent become the way in which apps are installed on an OS, it needs to allow competing channels (app stores, direct downloads) for users to get apps under Article 6(4) DMA. This obligation is clearly targeted at gatekeepers *with* an OS, and organizing app distribution through their own app store, though the obligation is formulated as to apply to any CPS. Hence, it could apply to an agent serving as app access channel whatever the CPS category it is designated in. Second, the obligation not to use data from business users to compete with them, found in Article 6(2) DMA, would also apply to agents designated in any category, even if it was adopted with online intermediation (online marketplaces) in mind. Third, there are the obligations concerning the choices available to business and end-users for communicating and transacting with each other – also outside of the gatekeeper platform – in Article 5(4) and 5(7) DMA. Again, these are generally applicable but kick in primarily when AI agents intermediate, e.g. complete purchases, in which case users should be able to choose how to pay for them.

In conclusion, when it comes to foreclosure by an AI agent, the in our view most fitting qualification as virtual assistant also leads to correct coverage in terms of obligations. When designing obligations, lawmakers seem to have viewed virtual assistants as (potential) middleware, i.e., a foundation for other apps or services, which is (coincidentally) in line with the expected development of AI agents. Hence, most of the obligations relating to OSs – and to some extent browsers – also apply to virtual assistants (e.g. changeable defaults/choice screens, equal vertical interoperability). AI agents also exhibit characteristics of search engines (and are in any case integrated with them), but the most important search-related obligations (in particular the prohibition of self-preferencing in ranking) is

¹⁴³ Even if most of them can also be put in terms of self-preferencing (of a gatekeeper's own app store or payment system, in terms of data access).



wide enough to also apply to virtual assistants. A final set of obligations applies independently of the CPS category. Hence, if AI agents are indeed designated as virtual assistants, the DMA does not immediately show gaps when it comes to expected concerns around foreclosure by AI agents. However, as such foreclosure scenarios are further in the future than those regarding foreclosure of AI agents, new gaps may appear over time. By contrast, other CPS qualification that are possible (but not preferred) in the short-term, i.e., the web browser or search engine categories, come with a less comprehensive set of obligations. Interoperability (e.g., of connected devices) would be a problem in the browser designation scenario, while the obligation of changeable defaults does not apply to search engines¹⁴⁴ – to give just two examples.

¹⁴⁴ Even if this could be compensated, at least in part, by the self-preferencing prohibition.



7. Conclusions and Recommendations

Despite its origin in earlier competition law cases, the DMA proves to be surprisingly well-suited to address the emerging era of AI agents. The DMA's adaptability to this context stems from two facts. First, the source of foreclosure of AI agents is likely to be found in CPSs that are already designated, particularly (mobile) OSs. Second, legislators added a "virtual assistant" category to the DMA's list of CPSs, and its definition can plausibly accommodate AI agents (by contrast, other CPS categories, such as browsers or search engines, offer a less appropriate fit when considering the expected development trajectories of AI agents and the specific obligations applicable under the DMA). Notwithstanding the relative suitability of the DMA to tackle potential issues regarding AI agents, some amendments could help secure its future-proofness. The amendments relate to the DMA's scope and obligations.

Regarding the DMA's scope, we recommend **mitigating potential ambiguities in the definition of "virtual assistant" (recommendation 1)**. Doing so raises two questions: one of substance (what to change); the other of institutions (how to change). On the substance, we discussed how the definition of "virtual assistant", which is built around the processing of prompts ("demands, tasks or questions"), captures AI agents in the near to mid-term. As AI agents attain ever more capabilities, autonomous action is likely to become more prevalent: the agents may anticipate user "needs" rather than reacting to prompts, which could thus be included in the definition.

As to how to mitigate these ambiguities, we discussed two routes: interpretation (via application) and legislative change. Each route has its benefits in terms of effectiveness, legal (and related investment) certainty, and political legitimacy. A good way to juggle these policy considerations is to view the routes as complements rather than substitutes. In the near to mid-term, the designation of AI agents in the virtual assistant CPS category is effective without significantly imperilling legal certainty; in the mid to long-term, legislative change provides a more legally certain way to capture the future iterations of AI agents.

Regarding the designation of gatekeepers, the existing quantitative thresholds appear sufficient at present. Companies such as OpenAI, with their early-stage AI agents, may already meet the user-based thresholds.¹⁴⁵ However, complications may arise in identifying and counting the business users of virtual assistants. Specifically, developers contacted by AI agents may not actively "use" these agents in the conventional sense. While some developers are taking steps to interoperate with AI agents, and more may do so in the future, others may instead be passively "called upon" by agents. This issue, however, could be resolved through interpretative adjustments aligned with existing practice for search engines. To increase legal certainty, the Commission could – via delegated act – adapt the DMA's Annex, including its definition of a virtual assistant's business users.

Given the rapid pace of technological change, the designation of "emerging gatekeepers" could be considered. In such cases, specific obligations under Articles 5(3)–(6) and 6(4), (7), (9), (10), and (13) DMA could be applied selectively. Nevertheless, the immediate priority remains the prevention of foreclosure of AI agents, for which the designation of OSs as gatekeepers is more critical than the

¹⁴⁵ For recent user numbers, see <https://techcrunch.com/2025/04/23/google-gemini-has-350m-monthly-users-reveals-court-hearing/>.



direct designation of AI agents themselves, except concerning user choice mechanisms under Article 6(3) DMA.

In situations where AI agents straddle CPS categories (e.g., displaying features of both browsers and search engines), or where current obligations are insufficient (such as in data-sharing requirements for search engines), multiple CPS designations might seem like an attractive solution. However, the DMA currently does not permit multiple designations, and any amendment to enable this would introduce significant legal and administrative complexity. Consequently, adapting the obligations attached to existing categories is a more feasible and efficient policy direction.

When it comes to obligations, we propose two adjustments so that the DMA adequately covers AI agents. The first adjustment is to **extend the click and query data-sharing obligations of Article 6(11) DMA to include virtual assistants (recommendation 2)**. Given the critical role of such data for AI agent contestability, this measure would enhance market openness. The second adjustment is to **add virtual assistants alongside browser engines and other services in Article 5(7) DMA (recommendation 3)**. This would help ensure that users retain freedom in their choice of AI agents and, as a result, foster competitive dynamics between such agents.

These proposed amendments aim to counter the foreclosure of AI agents. By contrast, concerns about foreclosure by AI agents do not currently warrant regulatory amendments, assuming AI agents are designated within the virtual assistant category (or in a future AI agent category).

A related concern pertains to the pre-installation of third-party AI agents. While Article 6(3) DMA addresses the pre-installation of a gatekeeper's own services, it does not explicitly cover contractual arrangements with third parties (e.g., partnerships between Samsung and Google, or Apple and OpenAI), in which there are signs of exclusion.¹⁴⁶ Although this issue lies somewhat outside the DMA's core logic – which primarily governs internal ecosystem dynamics rather than (contractual) inter-ecosystem relationships – emerging competition among third-party arrangements (e.g., between Motorola and Google, Perplexity and others) suggests that vigilance is warranted.¹⁴⁷

Importantly, the report distinguishes between concerns related to agentic frameworks and those involving foundation models, which we considered only as an input to AI agents. The DMA's provisions do not extend to such models nor their inputs, including training data and compute (also via hardware, such as AI chips). For example, if chip providers were to design their chips in such a way that they work better with certain AI models or agents, or embed specific models directly on-chip, then this would be outside the scope of the DMA. Further, the DMA does not ensure non-discriminatory access to off-device computing resources (e.g., cloud services). In this context, a possible side-recommendation is to **extend the CPSs covered by the FRAND access obligation of Article 6(12) DMA to include cloud computing services (recommendation 4)**. This could help ensure non-discriminatory access in an increasingly vertically integrated market.

¹⁴⁶ See <https://www.bloomberg.com/news/articles/2025-04-23/perplexity-executive-says-google-blocked-motorola-s-use-of-ai-assistant>.

¹⁴⁷ See <https://motorolanews.com/moto-ai-launches-new-experiences-and-partnerships-with-ai-leaders-giving-users-choice/>. However, the integration with Google's Gemini appears deeper (it is described as "deep integration"), while Perplexity appears to simply show up as an app on a phone, potentially without the required access to function as a capable agent.



More broadly, structural trends in the AI sector suggest that the **DMA alone may not suffice to prevent long-term concentration**. Vertical integration across the AI value chain – from chips and OSs to models and agentic frameworks – poses a significant challenge. Due to strong network and scale effects on both the demand and supply sides, increasing market concentration appears likely. This dynamic suggests that rather than seeking for competition *in* the market, policymakers should strive for competition *for* the market accompanied with the goal of transient monopolies, in line with Schumpeterian theories of creative destruction.¹⁴⁸

It is possible that today's gatekeepers – who control key resources like OSs and training data – will also dominate the agentic AI markets. New entrants such as OpenAI could themselves move upstream into hardware and operating systems, while incumbent players like NVIDIA and Qualcomm may integrate further downstream. However, strategic partnerships of existing gatekeepers with key players in the agentic AI value chain make it likely that the gatekeepers of today are also the gatekeepers of tomorrow, despite the DMA being a surprisingly good fit for the advent of agentic AI.

¹⁴⁸ Joseph Schumpeter, *Capitalism, Socialism and Democracy* (Routledge 2010 [1943]).



References

Autorité de la Concurrence, Opinion 24-A-05 of 28 June 2024 on the competitive functioning of the generative artificial intelligence sector, available at <https://www.autoritedelaconcurrence.fr/en/opinion/competitive-functioning-generative-artificial-intelligence-sector>.

Bostoen, F. and Krämer, J., 'AI Agents and Ecosystems Contestability' (CERRE Issue Paper, 2024).

Bostoen, F. and Petit, N., 'Antitrust rules and remedies against platforms' treacherous turns' (2024) 30 European Law Journal 629.

Competition & Markets Authority, 'AI Foundation Models' (Initial Report, September 2023), available at <https://www.gov.uk/government/publications/ai-foundation-models-initial-report>.

Competition & Markets Authority, 'AI Foundation Models' (Update Paper, April 2024), available at <https://www.gov.uk/government/publications/ai-foundation-models-update-paper>.

Draghi, M., 'The future of European competitiveness Part A | A competitiveness strategy for Europe' (Report, 2024), available at https://commission.europa.eu/topics/eu-competitiveness/draghi-report_en.

Draghi, M., 'The future of European competitiveness Part B | In-depth analysis and recommendations' (Report, 2024), available at https://commission.europa.eu/topics/eu-competitiveness/draghi-report_en.

European Commission, 'Sector Inquiry into Consumer Internet of Things' (Final Report) COM(2022)19 final.

European Commission, Impact Assessment Report accompanying the Proposal for a Regulation of the European Parliament and of the Council on contestable and fair markets in the digital sector (Staff Working Document) SWD(2020)363 final.

Graef, I. and Prüfer, J., 'Governance of data sharing: A law & economics proposal' (2021) 50 Research Policy.

Hagiu, A. and Wright, J., 'Data-enabled learning, network effects, and competitive advantage' (2023) 54 The RAND Journal of Economics 638.

Hagiu, A. and Wright, J., 'Artificial intelligence and competition policy' (2025) International Journal of Industrial Organization (advance access).

Hunt, S., Chissell, E. and Mawar, A., 'Will 2025 be the year of the agent? A primer for competition practitioners on the next wave of AI innovation' (2025) 9 Competition Law & Policy Debate 20.

Japan Fair Trade Commission, 'Generative AI and Competition' (Discussion Paper, October 2024), available at <https://www.jftc.go.jp/file/241002DiscussionPaperEN.pdf>.



Kak, A., Myers West, S. and Whittaker, M., 'Make no mistake—AI is owned by Big Tech' (*MIT Technology Review*, 5 December 2023) <https://www.technologyreview.com/2023/12/05/1084393/make-no-mistake-ai-is-owned-by-big-tech/>.

Kowalski, K, Volpin, C. and Zombori, Z., 'Competition in Generative AI and Virtual Worlds', European Commission Competition Policy Brief 3/2024.

Krämer, J., 'Personal Data Portability In The Platform Economy: Economic Implications And Policy Recommendations' (2021) 17 *Journal of Competition Law & Economics* 263.

Krämer, J. and Schnurr, D., 'Big data and digital markets contestability: Theory of harm and data access remedies' (2022) 18 *Journal of Competition Law & Economics* 255.

Krämer, J. and Shekhar, S., 'Regulating Digital Platform Ecosystems Through Data Sharing And Data Siloing: Consequences For Innovation And Welfare' (2025) 49 *MIS Quarterly* 123.

Malladi, D. and Chang, J., 'AI Disruption is driving innovation in on-device inference' (Qualcomm White Paper, 2025), available at <https://www.qualcomm.com/news/onq/2025/02/ai-disruption-is-driving-innovation-in-on-device-inference>.

Ménière, Y., 'Fair, Reasonable and Non-Discriminatory (FRAND) Licensing Terms: Research Analysis of a Controversial Concept' (JRC Science and Policy Report, 2015), available at <https://publications.jrc.ec.europa.eu/repository/bitstream/JRC96258/jrc96258.pdf>.

Meyers, Z. and Bourreau, M., 'A Competition Policy for Cloud and AI' (CERRE Issue Paper, 2025).

Morozovaite, V., 'The future of anticompetitive self-preferencing: analysis of hypernudging by voice assistants under article 102 TFEU' (2023) 19 *European Competition Journal* 410.

Noskova, V., 'Virtual assistants as gatekeepers for consumption? – how information intermediaries shape competition' (2023) 19 *European Competition Journal* 30.

Ofcom, 'Cloud services market study' (Final Report, October 2023).

Packer, C. et al., 'MemGPT: Towards LLMs as Operating Systems' (2024) arXiv:2310.08560.

Riley, C., 'Digging in on personal AI portability' (*Data Transfer Initiative*, 4 June 2024) <https://dtinit.org/blog/2024/06/04/digging-in-personal-AI>.

Riley, C., 'The future of AI hinges on data portability and APIs' (*Data Transfer Initiative*, 11 February 2025) <https://dtinit.org/blog/2025/02/11/future-of-AI-portability>

Schumpeter, J., *Capitalism, Socialism and Democracy* (Routledge 2010 [1943]).

Vasquez Duque, O., 'Taking Behavioral Antitrust Seriously: On Default Agreements as Exclusive Dealing and the Debiasing Potential of Default Randomization' (2024) 84 *Maryland Law Review* 143.

Ziegler, D. et al, 'Fine-tuning language models from human preferences' (2020) arXiv:1909.08593.

cerre

Centre on Regulation in Europe



Avenue Louise 475 (box 10)
1050 Brussels, Belgium
+32 2 230 83 60
info@cerre.eu
www.cerre.eu

 Centre on Regulation in Europe (CERRE)

 CERRE Think Tank

 CERRE Think Tank

