AI Foundation Models: Short Version

18 September 2023

AI Foundation Models: Short Version

1.1   Advances in machine learning and artificial intelligence (AI) continue apace. Recent developments in foundation models ('FMs') – large-scale models that can be adapted to a wide range of tasks and operations – and their rapid adoption across a growing number of user applications, have highlighted their potential to spur innovation and economic growth. These technologies are already being used to help us research, learn and solve problems.

1.2   We launched this initial review to help create an early understanding of:

- How the markets created or impacted by the development of FMs and their use could evolve, and the scenarios that may emerge;

- What opportunities and risks these scenarios could bring for competition and consumer protection; and

- Which principles can best guide the ongoing development of these markets so that the vibrant innovation that has characterised the current emerging phase is sustained, and the resulting benefits continue to flow for people, businesses and the economy.

1.3   The development of AI has raised several other important issues, including safety; security; privacy; intellectual property and copyright; and human rights. These issues are being considered by other regulators and Government. This review focused on questions that the CMA is mandated and best placed to address, namely questions around competition and consumer protection.

1.4   To inform our thinking, we have engaged with over 70 stakeholders, including a range of FM developers, businesses deploying FMs, consumer and industry organisations and academics. We gathered information directly from stakeholders as well as considering publicly available information, including the latest AI research.

1.5   Our analysis has focused on three levels of the value chain: (1) the development of FMs; (2) how FMs are used in other markets and user applications; and (3) the experience consumers have when using these new AI tools (either standalone or when incorporated in other products or services), in particular whether they can make informed choices and are treated fairly.

1.6   For each level in the value chain, we set out a broad spectrum covering the possible ways the market could develop from a competition and consumer protection standpoint, from positive outcomes[1] to outcomes that would cause concern. This approach provides a framework to consider the key drivers for either type of outcome and the development of a set of principles that can help guide the development of the market towards more positive outcomes for people, businesses and the economy.

1.7   Effective competition is crucial to ensure markets with healthy business rivalry, innovation and sustained productivity can flourish. Strong competition could spur the introduction of new products which help with all kinds of creative, scientific and administrative tasks, and benefit the whole economy by encouraging dynamism and growth. However, if competition is ineffective, the best firms, products or services will not necessarily win out, and both consumers and businesses may find they are locked into ecosystems with higher prices and restrictions that they cannot easily escape. We also recognise that effective competition alone is not sufficient to ensure good market outcomes. It is important to consider the role of effective competition alongside other considerations such as safety, data protection and intellectual property rights, for example.

1.8   This document sets out our early views on:

- How FMs are developed, the key inputs they require and how they are deployed today;

- The potential outcomes for competition in the development of FMs;

- The impact of FMs on competition in other markets and the potential outcomes for competition;

- The potential outcomes for consumers;

- The potential role for regulation in enabling positive development and outcomes;

---

[1] Given the inherent unpredictability of the future it would likely be impossible to create a comprehensive and accurate set of possible outcomes. We have not attempted to do so. Instead, these are necessarily stylised outcomes, and we do not claim that any of these options will materialise in the way we describe or at all. Rather they extrapolate, for analytic purposes, market features and trends that we think may emerge based on the evidence we have seen that we consider could have an impact on competition and consumers.

- Proposed competition and consumer protection principles that will guide the development of the market; and

- Next steps for the CMA.

How FMs are developed and deployed today

1.9 FMs are large, general machine learning models that are trained on vast amounts of data and can be adapted to a wide range of tasks and operations. Even focusing just on end-consumer applications, which are currently being used to power chatbots, create code writing assistants, and generate images and are being incorporated into some software, such as Microsoft 365, where they are helping users – both consumers and businesses – undertake tasks.

1.10 The first public FM – GPT – was released by OpenAI in 2018.[2] Since then, it is estimated that there are around 160 FMs that have been developed and released by a range of firms, including established players in other, already established markets, such as Google (which owns DeepMind), Meta, Microsoft and NVIDIA, as well as new AI companies such as OpenAI, Anthropic, Stability AI and Midjourney.[3] However, not all the FMs that have been created are currently in use or being monetised. As models develop in their capabilities, they can quickly become obsolete and replaced by newer models. For example, Google previously used its LaMDA family of models to power its Bard chatbot, but that has since been replaced by the more powerful PaLM-2 model.

Key inputs required for building a FM

1.11 Developing a FM requires ongoing access to:

- **Computing power** – FMs are large (many models have billions of parameters, trained on hundreds or thousands of gigabytes of data) and require significant computing power, both when they are trained and when they are used. Specialised chips used for training and running AI models are in very high demand relative to current supply. FM developers without their own computational resources typically enter an agreement or partnership with a cloud service provider (CSP).

---

[2] GPT was the first model based on transformer architecture. Subsequently, models based on this architecture have become known as foundation models ('FMs').
[3] Stanford University (2023) Ecosystem Graphs for Foundation Models.

- **Data** – FMs need vast quantities of training data to build the knowledge of the model (called *pre-training*) and when the models are refined for a specific application such as a customer service chatbot or a code writing assistant (an optional stage called *fine-tuning*).[4]

- **Technical expertise** – FM developers need highly skilled research scientists and engineers to develop and maintain competitive FMs.

- **Capital** – Building and maintaining a FM requires access to significant amounts of capital to fund the use of cloud services or supercomputers, a skilled workforce and possibly also the cost of high-quality data if it is not freely available.

1.12   To invest and compete effectively, FM developers need confidence that they can access these inputs.

How FMs are deployed and used in user-facing applications

1.13   As shown in Figure 1, once the models have been trained and fine-tuned, they can be released and deployed in user-facing applications in a range of ways, including by deploying FMs directly, accessing an externally controlled FM via APIs (ie 'AI-as-a-service'), or through building plug-ins that work with FM applications. These methods of deployment mean that a firm with a consumer-facing or a business customer business can incorporate FM technology into its business by using a developer's FM without the need to build and maintain its own FM.

---

[4] For completeness, some applications that use FMs, such as search engines, also require access to additional (often real-time) data to supply the necessary context or inputs when the model is used, eg to return output that makes use of current and relevant search results.
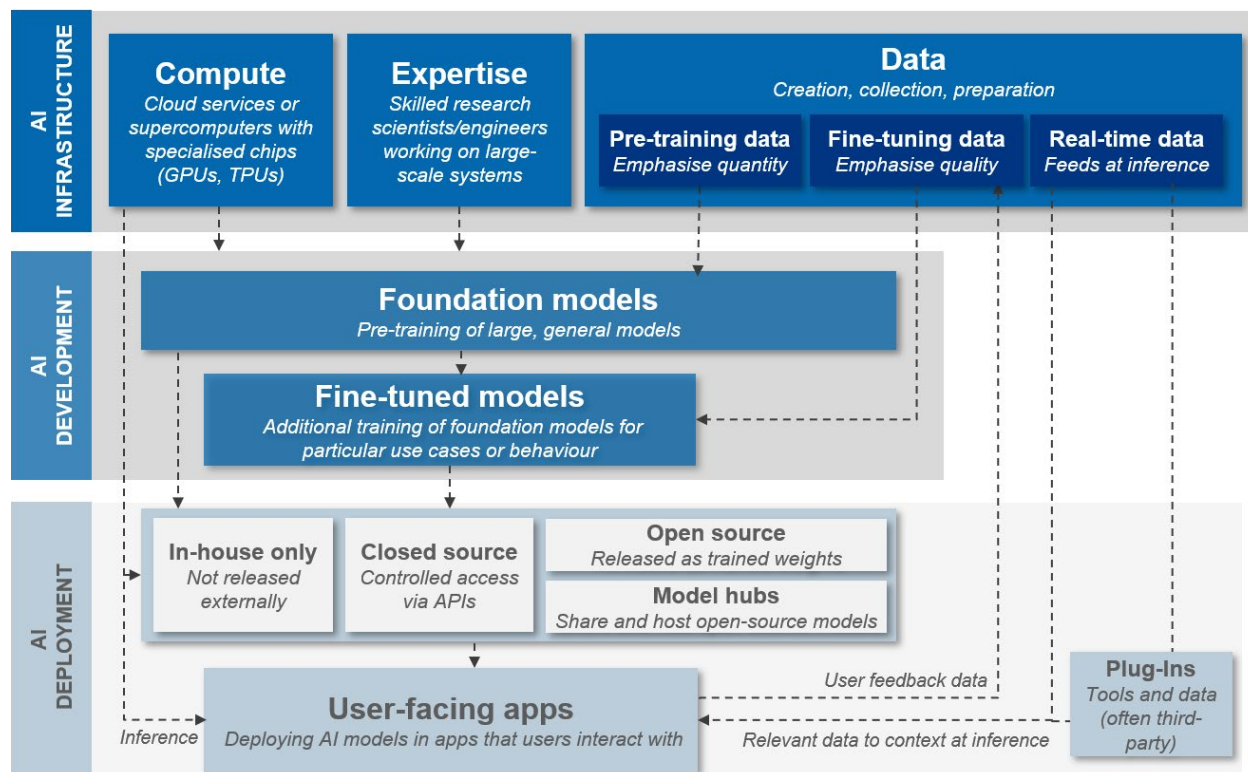
Figure 1: An overview of foundation model development, training and deployment

1.14　A key feature in this market is that approaches to release and making FMs available can vary in openness. **Open-source** models are those that have been released in a way which allows them to be freely shared for other FM developers to build upon to create their own FMs, and with relatively few or no restrictions on how they can be used (including for commercial use), such as the United Arab Emirates Technology Innovation Institute's Falcon model.[5] There are many models that are 'open' (in the sense that a copy of the trained model is readily available) but still have some licensing restrictions that limit commercial use, or restrict who is able to use it. For example, Meta's Llama-2 model is available freely for most commercial use, but if it is used in an app or service with more than 700 million monthly users, an additional license is required.

---

[5] There are a variety of ways in which FMs can be more open (including the availability of its code, data, weights, published information and documentation, and the permissiveness of its license), and that the term 'open' and 'open-source' are currently used in a variety of ways to describe FMs. In our report, unless otherwise stated, we have chosen to use 'open-source' in a way which emphasises the aspects of general availability of model weights and relatively permissive license to modify, extend and use for a variety of purposes (including commercial use), as these aspects are the most directly relevant to competitive dynamics.

1.15   In contrast, **closed-source** or proprietary models are not shared publicly, and there is often more limited public information about the models' characteristics and capabilities. Access and use of closed-source models are more controlled. Developers of closed source models can choose whether to use the models for only their own business or to license its use to other parties such as via APIs. For example, Bloomberg uses its FM called 'BloombergGPT' in its own financial services software 'Bloomberg Terminal' and does not license its use for others, and OpenAI controls access to its FM 'GPT-3' which it makes available via an API for third parties to use in their products and services.

1.16   At present a mix of open and closed-source FMs are available and competing. This is allowing a range of firms to invest in and develop FMs and as a result we are already seeing deployment of these FMs in a growing range of applications across the economy.

- **Search.** Microsoft has integrated models from OpenAI into its search engine Bing. Google has announced plans to incorporate FMs into search.[6] There are also many search or answer engines entering the market such as ChatGPT, You.com and Perplexity.ai.

- **Productivity software.** Google, Microsoft, Adobe, and Slack have all announced plans to integrate FMs into their existing products and environments.[7]

- **Social media**. Snapchat incorporated the ChatGPT-powered 'My AI' chatbot in its app that replies to users' posts or 'Snaps' with a text-based reply.[8]

- **Healthcare**. FMs are transforming scientific healthcare and drug discovery, including research on protein folding/expression prediction and rare disease research.

---

[6] Google - The Keyword (10/05/2023) How Google is improving Search with Generative AI; Microsoft Bing Blogs (2023): Confirmed: the new Bing runs on OpenAI's GPT-4

[7] Google Workspace (11/05/2023): Introducing Duet AI in Google Workspace; Microsoft (16/03/2023): Introducing Microsoft 365 Copilot ; Abode: AI art generator – Adobe Firefly; Slack: Introducing Slack GPT, the future of AI in Slack

[8] TechCrunch (31/03/2023) Snapchat launches a new generative AI feature, 'My AI Snaps,' for paid subscribers; The Verge (27/02/2023) Snapchat releases 'My AI' chatbot powered by ChatGPT

- **Robotics.** Researchers have been experimenting with FMs for a range of robotics applications including reasoning, planning, instructions and navigation.[9]

Firm structure and integration

1.17   Given the wide range of deployment options, a spectrum of possible firm and market structures may arise. Focusing on possible firm structures, one end of the spectrum could be a fully integrated firm which supplies its own computing power, develops its own FM using its own AI development tools, and deploys its FMs into its own products and services (integrations). At the other end, each stage in the value chain could be fulfilled by different firms (see Figure 3).

1.18   Today, we observe significant vertical integration, with many firms having a presence in two or more stages of the value chain. Several FM developers, such as Microsoft, Amazon and Google, own key infrastructure for producing and distributing FMs such as data centres, servers and data repositories.[10] Those firms also have a presence in a range of user-facing markets where FM technology can be integrated, from online shopping, search, through to the supply of software, so they have links across several parts of the value chain.


Figure 2 - A fully integrated value chain where the Cloud Service Provider provides all services.


Figure 3 - A non-integrated value chain where each service is provided by a different firm.

---

[9] For more information, see GitHub repository GT-RIPL/Awesome-LLM-Robotics for a list of papers experimenting with using FMs for robotics applications.

[10] J. Cobbe, M. Veale, J. Singh (2023) Understanding accountability in algorithmic supply chains.

1.19   We also see links across parts of the value chain in the form of partnerships and strategic investments. Google and Microsoft have entered into such agreements with various FM developers, including Anthropic[11] and OpenAI.[12] Both firms provide cloud computing services as part of their agreements.[13] We will closely monitor the impact of these investment, partnership and vertical integration links on competition in the development and use of FMs.

Competition in the development of FMs

1.20   To realise the full potential of FMs, it is vital that there is, on a sustained basis, effective competition between FM developers to produce high-quality models that can be deployed in a wide range of applications. A *positive market outcome* for people, businesses and the wider economy would arise if there were multiple independent developers competing with one another to produce leading FM models, with innovative firms able to access the inputs they need to enter, expand and compete effectively. In that scenario, firms would be able to experiment with different business models and forms of monetisation, including the supply of FMs on both an open-source and closed-source basis so others can continue to build on existing FM capabilities.

1.21   However, *a concerning market outcome* could emerge if access to inputs is restricted so only a handful of firms can create and maintain the leading models. As a result, those remaining firms would develop positions of strength which could give them the ability and incentive to provide models on a closed-source basis only and to impose unfair prices and terms. Any resulting reduction in competition may result in reduced incentives to innovate and this could reduce the scope for competitive innovation by a range of different firms, which may have a negative effect on economic growth and productivity.

---

[11] It is reported that Anthropic has received a total of $450 million in funding from Google. See here: www.reuters.com/markets/deals/alphabet-backed-ai-startup-anthropic-raises-450-million-funding-freeze-thaws-2023-05-23/

[12] Microsoft has invested a total of $13 billion in OpenAI over three rounds of funding. The first round, in July 2019, was for $1 billion. The second round, in January 2021, was for $1.5 billion. And the third round, in January 2023, was for $10 billion. See here: Microsoft to invest more in OpenAI as tech race heats up | Reuters

[13] Microsoft Corporate Blogs (2023) Microsoft and OpenAI extend partnership. Anthropic (2023) Anthropic Partners with Google Cloud.

1.22 Where on the spectrum of those two outcomes the market develops will be driven by:

- **Access to data** – the extent to which access to proprietary data will become necessary to compete effectively in developing FMs;

- **Requirements for and access to computing power** – the degree to which economically useful and relevant models will need to become larger and require more computing power (and other resources);

- **Whether large technology companies and first-movers have an advantage over others**; and

- **The existence of competitive open-source models** - will some competitive models remain available on an open-source basis, allowing FM developers to use and improve upon them without the need to build their own FM?

Access to data – Will access to proprietary data become necessary to compete?

1.23 Ready access to data has, to date, been a key factor in creating the conditions where new developers can experiment and develop new models, often developing models with comparable capabilities to the highest performing models. For example, popular FMs such as Meta's Llama 2 and Stability AI's Stable Diffusion were pre-trained using only data scraped from the web and other publicly available data.

1.24 However, we have heard that in future it could be more challenging for FM developers to improve on model performance by increasing the scale of training data because freely available data may be fully exploited (ie there is no new data that models could be trained upon) or grow at a slower rate. If that happens, to develop future FMs, developers may need to find ways to access new training data beyond what is freely available. Currently, FM developers have two options for sourcing new data: (1) they can use data they already have as a business, such as unpublished/private articles or analysis or (2) they can purchase data from third party providers, such as publishers and image repositories, in return for a fee and/or licencing conditions. For independent FM developers who do not already have access to relevant proprietary data, any shifts in the availability of data could affect the costs of developing an FM.

1.25   We have heard that some firms are already starting to use proprietary sources such as academic journals, image repositories, and content websites for model training, which suggests that the use of proprietary data is increasing in importance. If that trend continues, that could give an advantage to FM developers that already have access to good proprietary data. For example, a vertically integrated firm may gather useful feedback on how users interact with content on its social media platforms which can then be used to improve its FM to produce more relevant outputs or present results in a way users are more likely to engage with.

1.26   We considered whether synthetic data – data that is artificially generated (including by other FMs) rather than collected based on real-world events – could be used as a substitute to drive future improvements and provide FM developers with access to cheaper training data. We heard that there is a risk that the use of synthetic data generated by FMs for training future models could result in the irreversible degradation of their performance, a phenomenon referred to as 'model collapse'. There is ongoing research in this area, but we note that there is still considerable uncertainty about the extent to which synthetic data can be a viable data source for FM developers that could be a complete alternative to real-world and increasingly proprietary data.

1.27   There is therefore a risk that, if proprietary data becomes increasingly important to develop competitive FMs, but also less available and more expensive, many FM developers may have insufficient access to viable alternative data sources that they can use to keep pace. As a result, those FM developers may exit the market altogether or become reliant on a small number of firms to supply them with the necessary data. We cannot predict whether the market will tip entirely towards the use of proprietary data, but *the market will deliver better outcomes if it maintains a dynamic whereby a range of FM developers can gain access, on reasonable terms, to the data they need to build FMs. We would be concerned if firms used their leading positions in other markets to unduly restrict access to other competing FM developers.*

Requirements for and access to computing power – Will models need to become larger?

1.28   Although a detailed assessment of the supply of semiconductors was not the focus of this initial review, we understand that FMs require large, distributed computing systems, often consisting of hundreds of specialised chips used for training and running AI models, called AI accelerator chips. Currently, AI

accelerator chips are in very high demand relative to current supply, and they are expensive to acquire and have limited availability. NVIDIA is currently the main supplier of chips that are used for AI purposes, although other firms are in various stages of developing their own AI accelerator chips.

1.29   FMs have also been getting larger. One of the first FMs released was BERT in 2018 which had 354 million trainable parameters (values that encode the knowledge of the model).[14] Since then, models such as PaLM, GPT-3, and Megatron-Turing NLG have been developed with hundreds of billions of parameters,[15] and popular open-source models have in the range of tens of billions of parameters.[16]

1.30   The principal reason behind this trend is an observed positive relationship between scale and performance, known as 'scaling laws' – larger models, trained on more data, using more compute to train and run, tend to do better than smaller models. However, there is uncertainty over whether this relationship will endure if models continue to grow in the future or whether model performance could plateau or even decline at greater scale.

1.31   Larger models currently tend to do better but cost more to develop and use, particularly in relation to computing costs. Meta's FM 'LLaMA' has 65 billion parameters and an estimated compute cost of $4 million. In contrast, the larger FM 'Megatron-Turing NLG' with 530 billion parameters has an estimated compute cost of $100 million.[17] As a consequence, without significant investment, smaller FM developers are unlikely to be able to finance the computing costs required to train the largest models.

1.32   Pre-training FMs requires a large amount of computational power. Most FM developers do not own the sufficient computational infrastructure to train models in-house, therefore most rely on agreements or partnerships with CSPs. We have heard concerns that firms who already have agreements or partnerships with computing providers are more likely to get access to the computing power they need. Although some startups can receive investment in the form of 'credits'

[14] [1810.04805v2] BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding (arxiv.org)

[15] Pathways Language Model (PaLM): Scaling to 540 Billion Parameters for Breakthrough Performance – Google Research Blog (googleblog.com)

[2005.14165] Language Models are Few-Shot Learners (arxiv.org)

Megatron-LM: Training Multi-Billion Parameter Language Models Using Model Parallelism (arxiv.org)

[16] Open LLM Leaderboard - a Hugging Face Space by HuggingFaceH4

[17] Towards Data Science (2023) Estimating the Cost of Training LLMs | Towards Data Science, Hugging Face (2021) Large Language Models: A New Moore's Law? (huggingface.co).

from large computing providers to spend on cloud computing,[18] we have heard a concern that larger companies are still more likely to get 'first in line' and make deals to hold larger compute clusters.

1.33 Ofcom is currently conducting a market study of cloud infrastructure services in the UK. In its Interim Report, Ofcom highlighted cloud services as increasingly important inputs to many businesses and organisations across the economy, noting that cloud is also a cornerstone of recent technological innovations, including artificial intelligence.[19] Ofcom has provisionally identified features and practices that make it more difficult for customers to switch and use multiple cloud suppliers, and has proposed to refer public cloud infrastructure services to the CMA for further investigation.[20] Ofcom intends to publish a final report no later than 5 October 2023. In the event that Ofcom makes the market investigation reference, the CMA will carry out an independent investigation in relation to public cloud infrastructure services in the UK and determine whether there are any adverse effects on competition. This could include consideration of issues related to FM requirements and CSPs.

1.34 It remains to be seen how FMs will develop and how they will be adapted for different uses, and whether this will influence how large a model needs to be to perform tasks. Based on the evidence we have seen, it appears that, at present, smaller models do not offer the same level of performance as larger models, although this may change as AI technology develops. However, smaller models may nevertheless be an effective competitive option, as they can be developed and run more cheaply than larger models. Organisations will have different requirements for performance depending on their context and application. It is possible that some products and services will require cutting-edge performance and so require the highest performing FM, which is developed at a higher cost and only available from fewer providers. However, it is also possible that there could be a range of user applications that require good, but not cutting-edge, performance, where a smaller, cheaper, model would suffice to complete the task, with correspondingly lower barriers for providers and more options for customers. It is likely that over time a combination of differently sized FMs will be required, but it is unclear what the overall range of sizes might be, and how large the largest models might become.

---

[18] AWS Activate for Startups, Founders, & Entrepreneurs (amazon.com)
AI startup program | Google Cloud
[19] Consultation: Cloud services market study - Interim report (ofcom.org.uk), 5 April 2023, paragraphs 3.8-3.9.
[20] Ofcom proposes to refer UK cloud market for investigation - Ofcom

1.35   There is ongoing research into new FM architecture and design to identify whether there are more efficient ways to develop models that require fewer resources, like computing power. It remains unclear whether, as that research develops, models will become cutting-edge but on a smaller scale. However, if the trend towards ever larger models, with an increasing amount of computing power required to run them, continues it is possible that only the largest, well-resourced, firms or those with existing partnerships with computing firms will be able to develop competitive FMs. *The market is more likely to sustain a range of FMs, building models that are both at the frontier of new capabilities as well as performing routine tasks, if all FM developers can access computing power on fair commercial terms, without undue restrictions.*

Will large technology companies and first-movers have an advantage over others?

1.36   Firms may have an advantage in the development of FMs for a range of reasons. The most relevant potential advantages include:

- *Early mover advantages* – Early movers may have advantages in lower input costs and securing prominence and brand recognition. For example, learning effects – where early-movers and incumbents have longer and a larger customer base with which to drive and implement improvements – may give some firms the ability to convert an early lead into a durable advantage. It is uncertain how strong early mover advantages may be in practice or whether there will be an opportunity for late movers to catch up. It is also possible that late movers will have an advantage if they are able to free-ride or learn from the mistakes of early movers.

- *Need for funding and technical expertise* - The development of FMs requires significant funding and technical expertise. The cost of pre-training and fine-tuning large language models is high, and the process requires specialised skills and knowledge. Smaller FM developers that lack the necessary funding and expertise may be at a disadvantage in the development of FMs.

- *Economies of scale* – FM development may exhibit economies of scale, as initial high model development costs (pre-training, fine-tuning) can be spread over a larger customer base. There may also be more subtle effects where earlier FMs can be used in the development process of future models (eg code-assistant FMs can be used by developers creating FMs and FMs may

be used to generate data that can be used to train other FMs). It is uncertain how strong this factor may be, but if realised it could favour the technology companies which acquire an early lead in developing FMs.

- **Feedback loops** – where firms can use information they gain from a presence in other markets or from their wider customer base to improve their FMs – this could give those firms an advantage when improving their FM. For example, a vertically integrated firm that has a large volume of feedback from users (such as through thumbs-up or thumbs-down options in chat interfaces in consumer applications or through other user feedback data generated from licensed enterprise or productivity software) would have more real-life human feedback than a non-integrated FM developer to further enhance their products and services. The value of user feedback data for FM development is uncertain. Some stakeholders argue that, currently, once a FM completes its pre-training or fine-tuning, its performance level is essentially fixed, with the number of users having no immediate direct impact on user experience.

1.37 *Better outcomes may emerge in the longer term if new entrants can provide sufficient competitive constraints on the initially successful FM developers, despite potential advantages that these first movers may have from early movement, economies of scale and feedback loops. The market is also more likely to develop positively if firms with a leading position in other markets related to FMs do not use that position to foreclose rival FM developers and rivals in those other markets.*

The existence of open-source models: will some competitive models remain available on an open-source basis?

1.38 Open-source models allow a range of developers to access the underlying model weights, which developers can use to improve upon and build better models. We have seen that open-source access can spur innovation (see Box 1).

Box 1: An example of the use of open-source models

In March 2023, a pre-trained version of Meta's LLaMA model leaked online. Within 3 weeks, open research organisation LMSYS had fine-tuned LLaMA to develop a new FM 'Vicuna – 13B', using a data set of user-shared conversations with ChatGPT. LMSYS used GPT-4 to conduct 'a fun and non-scientific evaluation' of this new FM and estimated that the Vicuna-13B model achieved 'more than 90%*' of the quality of OpenAI's ChatGPT and Google's Bard – two of the most successful chatbots (see Vicuna: An Open-Source Chatbot Impressing GPT-4 with 90%* ChatGPT Quality | LMSYS Org).

1.39   There has been significant and rapid innovation occurring in open-source FMs, focusing on the development of high performing, finetuned models. Many developers are using publicly available pre-trained models instead of developing their own FMs, at a substantially reduced cost. For example, the Alpaca FM was developed by an academic institution for $600 based on a leak of Meta's LLaMA FM.[21] Open-source pre-trained models can therefore be a key part of helping reduce barriers to entry and ensuring that firms with innovative ideas can develop new models and improve existing ones.

1.40   We understand that there are several ways in which a model could be 'open', and to varying degrees. For example, a model release described as 'open-source' may consist of the underlying code, model architecture and training data so others can replicate the training process, but in other instances it can also include weights and biases (ie the 'knowledge') of the model, which allows others to use or fine-tune the model without conducting their own pre-training. Some researchers argue that the terms 'open' and 'open source' are used in a variety of ways, including for models that may not be open for others to scrutinise, use or fine-tune.[22] We also acknowledge that unconstrained use of open models could raise some safety concerns.

1.41   It is uncertain as to whether models will continue to be made available as open-source. In particular, the incentives to release or maintain models on an open-source basis are likely to be affected by a range of factors including the ongoing cost of computing power and how FMs are monetised in future. Whilst we would not expect all models to be open-source, *the market for FMs is more likely to trend towards positive competitive outcomes if there are a range of models pushing at the frontier of new capabilities. These would include models made available on both an open-source and closed-source basis. This would allow a broad range of actors to enter the market and drive innovation and give customers a wide range of choices of FMs.*

Conclusion

1.42   The development and competitiveness of FMs depend on access to key inputs, such as computing power, data, expertise, and funding. If these key inputs

---

[21] Whilst Meta's first LLaMA model was not licensed for commercial use, there are other pre-trained models such as Falcon-40B or LLaMA-2 that are, and could similarly be fine-tuned by firms at a lower cost than developing pre-trained models from scratch. Stanford CRFM

[22] See, for instance, Widder, D. G., West, S., & Whittaker, M. (2023), Open (For Business): Big Tech, Concentrated Power, and the Political Economy of Open AI.

become constrained, smaller FM developers may find it challenging to compete with larger, well-established businesses that have greater resources. This could lead to a decline in competition and innovation within the FM sector, which could ultimately harm consumers.

1.43 We have also identified a number of key uncertainties regarding the future development of FMs. The impact of these uncertainties on competition in the FM sector is not yet known, but they could make it harder for some firms to compete effectively, which could stifle innovation and limit diversity in approaches:

- The need for proprietary data for training FMs, which could create a barrier to entry for smaller organisations.

- The need for larger models, which could disadvantage smaller organisations with limited resources.

- The need for cutting-edge performance, which could place a significant burden on smaller organisations without the required resources or expertise.

- The advantages that large technology companies may have due to their access to vast amounts of funding, data, technical expertise, and resources.

- The challenges facing open-source FM models, including possible licensing restrictions, funding uncertainty, and the potential for closed-source models to outperform open-source models in the longer-term.

Given the likely importance of FMs across the economy, we would be concerned if access to the key inputs required to develop FMs were unduly restricted, in particular restrictions on data or computing power. The market is more likely to trend towards positive outcomes if:

- A range of FM developers can access the key inputs they need to build FMs, including data, computing power, capital and expertise, on fair commercial terms without undue restrictions.

- Initial successful FM developers face an ongoing competitive constraint from new entrants, so they do not gain an entrenched and disproportionate advantage by being an early mover in the market, having economies of scale or benefiting from feedback loops.

- There are a range of models – including open-source and closed-source models and FMs pushing at the frontier of new capabilities – available for firms to choose from.

- Firms are unable to use their leading position in other markets to unduly restrict access to firms they compete with in those markets or other competing FM developers.

The impact of FMs on competition in other markets

1.44   FMs are used in a variety of applications across a wide range of industries, from healthcare, finance, education, music generation, legal and many more (we refer to these as 'downstream FM services'). As FMs become more widely used, it will be important to ensure effective competition both in the development of FMs and in industries applying FMs. Vibrant competition will encourage innovation, ensure that the best products are successful and lead to the lowest possible prices for consumers. On the other hand, if competition is weak consumers may be forced to use worse products at higher prices.

1.45   FM deployment is still at an early stage, and the eventual range of outcomes is still uncertain, but we can characterise two broad ways FMs can be used:

- **Improving existing products and services** – Firms are incorporating FMs into their existing processes, products and services, either at the user-facing level or earlier on in the value chain. This could allow firms of many shapes and sizes to be more productive and efficient. In some markets, this can make it easier for a wider range of businesses to develop products and services. Examples include creative industries, such as marketing, where FMs are being used by firms – large and small - to produce materials, such as visuals, which may allow them to compete more effectively with market incumbents. FMs are also being used to improve existing services, such as providing new functionality, in markets such as online search and productivity software (see case studies below). Such improvements to products and services can be good for users. In some situations, they can, in principle, also affect competitive forces by either entrenching or undermining existing strong market positions. For example, in productivity software, FM services are being incorporated into products and services where firms have existing strong positions eg incorporating an AI powered virtual assistant into word processing software (see case study Box 3 on 'Productivity Software'). Incorporating FMs could, therefore, entrench the strong positions of incumbents. On the other hand, a number of start-ups are also offering FM-powered productivity software services that have the potential to disrupt incumbency positions.

- **Creating new products and services** – Future markets are inherently uncertain, but we can anticipate that FMs may be used to build new types of services that offer new solutions for people and businesses. Some of these could further reinforce the positions of incumbents in adjacent or linked

markets, while others could undermine existing strong market positions. For example, in search, there has been speculation that new services like ChatGPT could challenge Google's position (see case study Box 2 on 'Online search'). An example of a service that could emerge, but not yet seen in full form, is a generally capable virtual personal assistant, which incorporates FM and non-FM services to offer consumers support across a whole range of needs. It is, of course, impossible to accurately assess what the impact on competition will be from potential new FM products and services. As with many new ideas, such products and services have the potential to disrupt markets and existing market power. On the other hand, they could also create new or entrench existing positions of market power for the firms that develop that product or service. Where such products and services are developed by multiple firms who can compete effectively, this is less likely. On the other hand, where the product or service is developed by a single (or small number of firms) because they are able to leverage advantages and market power they hold in related markets (such as, unique or privileged access to proprietary data, entrenched distribution, virtual lock in of customers into an integrated ecosystem of services, etc), this may make it more likely that the market will become concentrated with only a small number of players able to acquire market power in these new sectors and/or to lock customers into broader ecosystems of related products and services from that firm.

1.46   We have heard that many firms are still considering how FMs could be used in their products and services. It is therefore too early to say how important FMs (and FM-enabled features) will become as an input into different downstream FM services. However, irrespective of the way that firms deploy FMs in their businesses, the market is more likely to produce *positive outcomes* if:

- downstream firms have access to a range of FMs and can switch between them easily, without being encumbered by unfair restrictions;

- users have access to a range of FM services and are able to easily switch and make active and informed choices about the best FM product and service for their needs; and

- competition to develop downstream FM services is not constrained by anti-competitive provisions in agreements or anti-competitive conduct by firms with market power (eg a vertically integrated firm foreclosing upstream or downstream rivals, terms in a partnership agreement that restrict the other firm's ability to compete effectively, or anti-competitive bundling or tying).

1.47   However, *a more concerning market outcome* could arise where:

- there is a lack of competition to develop FMs and/or downstream FM services, leading to a lack of choice for downstream customers eg because firms with market power in adjacent or downstream markets are able to leverage that market power to restrict competition in FM development;

- downstream customers have difficulty switching between FMs or FM products and services eg because they are locked into ecosystems that only offer a limited range of FM deployment options or FM products and services;

- firms who acquire market power in FMs, use this power to restrict competition in adjacent or downstream markets eg by foreclosing access to downstream rivals; and/or

- FMs are used to entrench market power in downstream or adjacent markets, potentially allowing firms to leverage that market power to unfairly disadvantage rivals and reduce competition in those markets or related markets eg through anti-competitive tying or bundling of FM products and services.

1.48   The likelihood of a positive or concerning outcome will be driven by a range of factors including:

- **Effective choice and the ability to switch** – will firms have access to a range of options to incorporate FMs into their services and will they find it easy to switch between them?

- **Customer preferences** – will customers prefer FM services offered within integrated systems?

- **The impact of vertical integration** – will integrated firms and those with partnership relationships have an incentive to foreclose upstream and downstream competitors?

- **Feedback effects** – how significant will data feedback effects be?

Effective choice and the ability to switch

1.49   Downstream firms can deploy FMs to improve existing products and services or develop new ones. Firms have a range of options available to them, they can:

20

- **Develop a FM in-house from scratch** – where the firm takes responsibility for creating and maintaining the FM and applying it to their own products and services. Firms that adopt this approach include Bloomberg, Pfizer, and Adobe.[23]

- **Partner with a FM provider to enhance an existing FM** – the firm may use a third-party FM, which removes the need to develop the model itself, but fine-tune the model with its own proprietary data to tailor it to its business needs.

- **Buy API access to a third-party FM and FM deployment tools[24]** - the firm could buy API access to a FM and FM deployment tools. Firms that currently adopt this approach include Duolingo, Shutterstock and Expedia.[25]

- **Use a third-party plug-in** – A firm can develop a plug-in to augment its offering with an FM-based service, such as ChatGPT.

1.50   Stakeholders, including FM service developers, told us that they currently find it relatively easy to switch between options. Being able to switch easily may help ensure that that competition between FM developers is intense as they need to work hard to retain their business customers.

1.51   It is uncertain whether it will remain easy and affordable to switch between alternative options for deploying FMs over time. This will depend in part on whether the market for the development of FMs is competitive, which will drive the availability of a range of models that firms can choose from. Competition at the FM development level in the value chain will, in part, be more effective if FM service developers are able to easily switch between FMs. *The market will continue to support switching and multi-homing, if firms (a) have a range of FMs options and systems to choose from when deciding how to adopt FMs in their businesses and (b) if FMs and the systems they use are*

---

[23] Introducing BloombergGPT, Bloomberg's 50-billion parameter large language model, purpose-built from scratch for finance | Press | Bloomberg LP ; Pfizer Doubles Down on AI/ML to Bring Transformative Medicines to Patients | BioSpace ; Bringing Generative AI into Creative Cloud with Adobe Firefly | Adobe Blog

[24] See emerging market for pre-trained FMs and FM deployment tools (eg ChatGPT API, Google Bard API (beta version), NVIDIA AI Foundations, IBM Watson x).

[25] SHUTTERSTOCK PARTNERS WITH OPENAI AND LEADS THE WAY TO BRING AI-GENERATED CONTENT TO ALL - Press and Media - Shutterstock ; Expedia Group ; Duolingo Max Uses OpenAI's GPT-4 For New Learning Features.

interoperable with one another. This will help ensure firms are not locked into one ecosystem.

Customer preferences

1.52   Consumers and businesses currently access downstream FM services in several ways. Some FM services, such as chatbots, may require consumers to visit a specific website or app to access them. Others, like writing assistants, may come bundled with specific products and must be accessed through those products.[26] As discussed, above, some customers may prefer to use a platform which gives them access to a range of FMs that they can use depending on their business needs. Consumers may also value the convenience of being able to access an ecosystem of FM services and non-FM services at once. Indeed, some firms are already integrating their services in this way, for example in productivity software and operating systems.

1.53   Users may be attracted to the convenience of having a single, integrated ecosystem, and this may deliver benefits for them. For example, an ecosystem could learn about a user's writing style from how a user engages with other services, such as mobile, search and email, and use that to create a more customised word processing software. This customisation will likely improve the customer experience by making the software more intuitive to use, resulting in efficiency and productivity benefits. In turn, the engagement with the word processing software can then be used to further customise the other tools in the ecosystem as well as enhancing the convenience and ease of use of the ecosystem itself by further adapting it to the user's requirements, with further resulting efficiency and productivity benefits.

1.54   Customisation will improve the customer's experience of FM services, but unless they can also move to a rival provider without losing the benefits of customisation that may be built over years, they will be reluctant or unable to switch FM provider because they will need to 'start from scratch'. As a result, over time, customers – both businesses and consumers – may find that they are locked into a single ecosystem thereby limiting their freedom to choose innovative standalone services, which will weaken competition. The extent of consumer 'lock-in' will partly depend on whether data portability is technically feasible across different FM service ecosystems, so that the new service can replicate

---

[26] For example, GitHub Copilot, a FM-powered code editor, is built into and accessed through GitHub, a platform for collaborative software development. GitHub Copilot · Your AI pair programmer · GitHub

some or most of the value from customisation of the previous service.[27] *The market is likely to avoid the risk of consumers being locked in to closed ecosystems if they are able to port their data easily between services so they do not have to 'start from scratch' when switching between services and/or using multiple FM services at the same time ('multi-homing').* This outcome is likely to foster the greatest amount of competitive innovation.

The impact of vertical integration and partnerships

1.55   As discussed above, there is already significant vertical integration with many firms having a presence in two or more stages of the FM value chain and a number of partnerships between FM developers. There can be benefits to vertical integration – efficiencies, enhanced capabilities, and the ability for users to access a wider range of services that work well together. Partnerships can also help drive competition by supporting smaller firms that could not otherwise compete in the upstream FM development and supply and/or in downstream FM service markets.

1.56   In some cases, an integrated approach allows firms to provide customers with more choice. For example, Amazon Bedrock – a platform service that gives customers access to FMs - provides users with access to Amazon's own FMs as well models from other suppliers such as Anthropic and Stability AI.[28] Such integrated platforms may have benefits to businesses and consumers by offering convenient access to a range of models.

1.57   However, vertical integration and partnerships can also create opportunities for firms to distort competition in markets which could give rise to less positive market outcomes. For example, we would be concerned if a vertically integrated firm with market power imposed restrictive terms in a partnership agreement that prevented that other firm from competing effectively with it in downstream markets.

---

[27] Some FM services today allow consumers to export their data for external back-up or use (For example, as of April 2023 ChatGPT history can be exported and downloaded (https://help.openai.com/en/articles/7260999-how-do-i-export-my-chatgpt-history-and-data)). But it is unknown whether that data is compatible across different ecosystems.
[28] AWS announces Amazon Bedrock and multiple generative AI services and capabilities (aboutamazon.co.uk)

Box 2: Case study on online search

<div style="border: 1px solid blue; padding: 1em;">

Case study: Online search

Some stakeholders told us that FMs will be important for competing in search. The adoption of FMs by leading search engine providers could strengthen their positions in online search, as they may be best placed to develop and implement these new technologies effectively.

In addition, search engine providers may have some competitive advantages over other FM providers. For example, to provide general search, search engine providers already collect and process web crawl data, which may be more extensive and higher quality (after filtering and processing to create a search index) than publicly available web crawl datasets and can serve as training data for FMs.

Some stakeholders said that current barriers to entry such as control over default settings continue to give incumbents a significant advantage. On the other hand, FMs can benefit smaller search engines by enabling them to leverage advanced language capabilities without extensive resources. This could help them compete with larger players in online search.

The potential impact of FMs on online search is complex and will depend on various factors. While some commentators have suggested that FMs could strengthen current positions in search, as above, others have suggested that chatbots or 'answer engines' could supplant search engines. It remains uncertain which outcome is more likely at this stage.

</div>

Box 3: Case study on productivity software

Case study: Productivity software

Google, Microsoft, Adobe, and Slack are all integrating FM-based features into their existing productivity software services. Firms can direct their existing consumer base to new FM-based features and use data from their existing services to develop those features.

Microsoft is testing a Copilot virtual assistant that allows users to seamlessly access Outlook email, Office suite, Bing search and Azure cloud services as well as a range of third-party plug-in providers. The Copilot can draw on materials across these services, eg create a presentation combining data from a Bing search and the user's own Office-created documents.

Google recently announced an experimental product, NotebookLM, that is an FM-powered 'virtual research assistant' that can summarise facts, explain complex ideas and help brainstorm, all personalised to the consumer's own notes and sources.

In future, it is possible that virtual assistants like NotebookLM could give rise to customised, integrated ecosystems of FM-powered productivity software services and many other FM services. This could be beneficial to consumers but could also make it harder for firms offering standalone FM services to compete.

The potential impact of FMs on productivity software is unclear and will depend on factors such as how FM-powered features are monetised, the extent to which consumers value integrated and/or customised ecosystems and how easily they can switch between them.

1.58 Vertically integrated firms may also have the ability and incentive to withhold or limit access to key inputs or routes to market (exclusionary strategies) to favour their own downstream services or extract an excessive proportion of the value generated by rivals (exploitative strategies). Vertically integrated firms with upstream market power may be incentivised to foreclose downstream rivals in cases where there is greater potential for monetisation downstream than upstream.

1.59 A positive outcome would be one where the benefits of vertical integration and partnerships can be realised, without harming competition or restricting businesses' choices. To ensure the market trends towards such an outcome, *it is important that businesses are not subject to anti-competitive conduct. This includes anti-competitive self-preferencing, tying or bundling. It is important that businesses are able to invest in FMs and FM services with the confidence that they can compete on the merits in these markets.*

How significant will data feedback effects be?

1.60 In the context of AI, 'data feedback effects' refer to the ability of FMs and FM developers to use data generated by their usage to improve their performance. The extent of data feedback effects will be an important determinant of how competition in the downstream market for FM services will develop. Generally, the greater the feedback effects, the quicker firms will be able to make their downstream FM services better, giving these firms a competitive advantage.

1.61 User data can improve FM performance and this data can be gathered in different ways. Firms can gather data from FM services through downstream user feedback (such as a 'thumbs up' or 'thumbs down' to a specific output), or from how the consumer interacts with a service (eg what questions they ask, and how they react to the response). This data can be used to either customise the downstream FM service to consumer preferences or fine-tune to improve the FM's general performance, potentially strengthening the firm's position when competing with other FM developers.

1.62 It is too early to draw conclusions about the significance of data feedback effects. A key factor determining the extent of the likely effects will be whether the data that firms gather from downstream services can be fed into FMs in real time to improve their capabilities. If that occurs, it could accentuate the existing impact of data feedback effects and create 'first mover' advantages that make downstream markets more likely to tip towards concentration and make it more difficult for rival FM services developers to compete*. The market is more likely to develop positively if markets are open and competitive where FM developers and deployers are subject to competitive constraints which weaken the effect of any possible advantages that may emerge in the future, such as data feedback effects or first mover advantages.*

Conclusion

1.63   The evidence points towards the potential for FMs to transform a wide range of services and be the catalyst for significant innovation and competitive disruption in markets. However, it is unclear exactly how important FMs will be, and which downstream markets are most likely to be affected. One of the most important factors in ensuring FMs drive competition in downstream markets is the continued availability of a wide range of innovative deployment options that downstream firms can easily switch between.

1.64   A competitive upstream market allows downstream firms to switch between FM models. That, to a certain extent, helps mitigate against possible harmful effects in downstream markets because downstream firms would have a choice between FMs and could easily switch and multi-home. However, it does not eradicate the potential for harm entirely. Even with competitive conditions upstream – where FMs are developed by a range of different firms who compete vigorously with one another to improve their products – factors in the downstream market could still cause harm. These factors include: (1) significant data feedback effects that may tip the downstream market towards concentration (2) consumers' inability to meaningfully choose between FM services or switch away from downstream ecosystems (3) vertical integration and partnerships that harm competition or restrict businesses' choices.

> The market is more likely to produce positive outcomes if:
>
> - Firms can choose between a range of options when deciding how to adopt FMs in their businesses.
>
> - FMs and the systems they use are interoperable with one another.
>
> - Consumers can port their data easily between services, so they do not have to 'start from scratch' when wanting to switch or use multiple FM services.
>
> - Businesses are not subject to anti-competitive conduct, including anti-competitive self-preferencing, tying or bundling.
>
> - The market is more likely to develop positively if markets are open and competitive where FM developers and deployers are subject to competitive constraints which weaken the effect of any possible advantages that may emerge in the future, such as data feedback effects or first mover advantages.

1.65   We would be particularly concerned if we saw firms unfairly gaining or entrenching their market positions through leveraging their positions in adjacent downstream markets or in the upstream development of FMs, including as providers of key inputs to FMs.

Consumer protection

1.66   New uses of AI have the potential to enhance consumers' experience of products and services. However, for consumers to make the most of what new AI services have to offer FM developers and businesses that use FM-based services need to ensure that the outputs consumers receive are reliable and accurate and they must treat consumers fairly. This is particularly important because, as discussed, FMs and new AI technology is complex, and consumers may not always be aware that they are using an FM-based service.

1.67   The development of FMs raises a number of important issues, for example relating to safety, security, copyright, privacy and human rights, which government or other authorities are better placed to address. As the UK's primary competition and consumer authority, we focused on how FM services may harm consumers if they give them false or misleading information (called 'hallucinations') or manipulate them in other ways. A survey by Deloitte estimates that more than 4 million people in the UK have used FM services already,[29] showing that consumers are already making use of new AI services. However, it also means that the impact of any inaccurate or misleading content on consumers could be significant, so it is vital that consumers are sufficiently protected.

1.68   A *positive outcome* would arise if FM-outputs were accurate, reliable and in the consumer's interest so they could have confidence in using a range of AI enabled products and services. However, *a concerning outcome* could persist where FMs continue to 'hallucinate' and provide users with inaccurate or misleading results or where FMs are deliberately used to help trick or mislead consumers. Consumers may also receive results that are not in their interest, for example, they could be given chatbot answers that are driven by commercial incentives rather than the best options for their needs.

---

[29] Deloitte (2023) More than four million people in the UK have used Generative AI for work - Deloitte

1.69 The key drivers for whether the market develops in a positive way for consumer protection include:

- **Improvements in model accuracy** – will FMs become more accurate over time and less likely to 'hallucinate'?

- **Clear accountability** – will it be clear who is responsible for FM outputs and how consumers can get redress if things go wrong?

- **Providing the right information so people can make informed choices**– will businesses and consumers receive the right information about FM products and services so they can make informed choices?

- **The ability of consumers to understand and critically assess FM outputs** – will consumers know when they are interacting with AI and how best to use FM tools and understand both their benefits and limitations?

- **Action to tackle consumer harms** – to what extent will FM tools exacerbate existing consumer harms?

Will model accuracy improve?

1.70 Currently, FM-generated outputs can be false and misleading, due to the tendency for the models to 'hallucinate' – to produce convincing, but incorrect outputs. This issue has deep technical root causes and is not easily solved. Although some progress is being made, we heard during our review that it is unlikely that we will reach a point where hallucinations are eradicated altogether. If that is the case, businesses will need to carefully consider the impact of any false and misleading information on consumers' decision making and take steps to ensure they are sufficiently protected.

1.71 *The market is more likely to trend towards a positive outcome where models are as accurate as they can be if FM developers face competitive pressure to improve the reliability and accuracy of their models and all firms take steps to ensure they fully comply with consumer law.*

Will consumers have clear routes to redress if things go wrong?

1.72 In complex sectors, where FMs are produced, deployed and used within a supply chain by a range of actors, it can be difficult for consumers to identify who is accountable and responsible if something goes wrong. If firms are not held

properly accountable for their role, they may have reduced incentives to invest in strategies to address consumer harms, such as by investing to reduce instances of hallucinations.

1.73   *The market will produce better protection for consumers if there is a mechanism to determine the proper allocation of accountability and responsibility*, so all firms have the right incentives to improve consumers' experience and they can seek redress when things go wrong.

Will consumers be able to understand and critically assess FM outputs?

1.74   The integration of FMs into consumer facing applications is relatively new and therefore there is limited research and analysis in this area. However, the limited research that we have seen, along with concerns expressed by some stakeholders during our review, suggest that people generally find it difficult to tell when they are interacting with FM content (unless this is clearly labelled) and the difference between information that is faithful and factual or not when it is FM-generated. This has the potential to lead to consumer harm if consumers are making decisions based on information that is false or misleading or a mistaken understanding as to the source of the information. It is likely that consumer understanding will improve over time, but we have not yet seen evidence to demonstrate that current understanding is developed enough to overcome the risk that consumers will not be able to tell the difference between FM and human-generated information and will act upon false or misleading FM-generated information.

1.75   Stakeholders also raised with us the issue of possible manipulation of users by FMs. Without measures to improve safety, FMs could conceivably engage in deceptive or manipulative conduct.

1.76   *Consumers are more likely to make informed choices that meet their needs and interests if they have the right information, including whether content is FM-generated, and the risks and limitations associated with that FM and any content it produces.*

Will there be reliable information available that businesses and consumers can access, assess and act on?

1.77   For businesses and consumers to drive effective competition on dimensions like accuracy, reliability and safety, it is vital that they can make informed choices,

including through understanding the performance and limitations of FM products and services. Consumers need to be able to compare and switch between products and services that use AI. Equally, for businesses to incorporate AI into their products and services, to be able to satisfy their own obligations under consumer law, they also need reliable information on the relevant characteristics of models[30] to enable them to manage their own risks and prevent harm to consumers.

1.78 *The market is more likely to develop positively if FM developers provide sufficient, understandable and accurate information to their customers (both business and consumers) to enable them to make informed choices and to ensure that consumers are protected from harm.*

Could FM tools exacerbate existing consumer harms and how can that be prevented?

1.79 It is uncertain how far developments in model technology will go in eradicating false and misleading content, but today we see that there is a significant risk that AI outputs could potentially exacerbate existing harms, including:

- **Fake reviews** – Fake and misleading reviews for products and services can lead to people making poorly informed choices and/or buying the wrong products and services. The increased use of FM tools may in future make it easier and cheaper for bad actors to create fake reviews. Moreover, it can be difficult to tell the difference between a genuine and a fake review, and FMs may make that problem worse because they could generate content that may be even more convincing. It is unclear whether FM tools will be used in this manner or what effect they may have. For example, it is also possible that FM tools will help firms to better identify fake reviews where they arise.

- **Phishing** – Tactics deployed by criminals to convince consumers through scam emails, texts or phone calls to disclose personal information or make payments[31] could be exacerbated if FMs are able to produce even more convincing, personalised content at scale. Research has found that FMs can

---

[30] Where access to FMs is provided as an ongoing service, customers may also need access to appropriate and relevant key information about FM providers' processes, such as how models are developed and maintained, quality management and other governance processes.

[31] Phishing: Spot and report scam emails, texts, websites and... - NCSC.GOV.UK

generate phishing emails that are difficult to detect and that have a high success rate in tricking individuals.[32]

- **Deep fakes –** Where FM-generated photos and videos are produced depicting fictional events, sometimes including generated images or voices of real people. Studies show that it can be difficult for people to tell the difference between real and deepfake content.[33] Despite that, one study shows that people overestimate their abilities to detect deepfakes – people detected deepfakes 42-77% of the time but reported confidence in their judgement 73-85% of the time.[34]

- **Hidden advertising –** Hidden advertising is harmful and illegal. The CMA has issued guidance relating to hidden advertising online which highlighted, among other things, the importance of ensuring that advertising and other commercial content is clearly recognisable as soon as a consumer engages with it[35] and that it is clear to consumers when the response they receive to a query[36] is affected by the money a business earns.[37] Although we are not aware of any FM providers currently using advertising as an input or prompt for their FMs that generate content or answers for users, there may be scope for this to happen in the future. There may be a breach of consumer law where commercial content produced by an FM to a consumer contains hidden product or service advertising.

1.80   *If the market is to mitigate some of these possible outcomes, it is important that sufficient safeguards exist to protect consumers from bad actors using FMs.*

1.81   Businesses that incorporate FMs in their consumer facing products or services should consider carefully whether they have satisfied their obligations under consumer law[38] and they should also keep this under review as practices, technology and the law continue to develop.

---

[32] Karanjai, R, (2023), Targeted Phishing Campaigns using Large Scale Language Models

[33] Bray, Sergi & Johnson, Shane & Kleinberg, Bennett, (2023), Testing Human Ability To Detect "Deepfake" Images of Human Faces. In fact, models may be better than humans at detecting digital forgeries – see Rössler, A. & Cozzolino, D & Verdoliva, L. & Riess, C. & Thies, J. & Nießner, M., (2019), FaceForensics++: Learning to Detect Manipulated Facial Images.

[34] Bray et al. (2022) Testing Human Ability To Detect Deepfake Images of Human Faces (arxiv.org)

[35] See CMA Guidance on Hidden ads: Principles for social media platforms, 3 November 2022.

[36] This includes, for example, product rankings and 'premium' listings.

[37] See CMA Blog on Accommodation booking sites: how to comply with consumer law, 26 November 2019.

[38] How consumer law applies is set out in chapter 6 of the main report.

Conclusion

1.82 The adoption of FM services is likely to be *more positive if consumers are sufficiently protected from the harms that may result from the use of FM tools by them or others and if both consumers and businesses are fully informed of their risks and limitations.* To achieve this, FM developers must develop models that are at the cutting-edge of both reliability and accuracy for consumers, as well as performance. In addition, consumers need accurate and reliable information about the products and services they are using to make informed and effective decisions.

1.83 We have heard about various efforts and strategies to make FMs more trustworthy, including technical measures, governance, disclosures, and consumer education. (For example, in July 2023, Anthropic, Google, Microsoft and Open AI announced the launch of the 'Frontier Model Forum ('FMF')', an industry body focused on ensuring safe and responsible development of frontier AI models.[39]) These are encouraging and we welcome robust action going forward to make FMs safer and more reliable for consumers.

1.84 Firms that use FM-services in their own products and services, however, should be mindful of their obligations under consumer law to ensure consumers are not harmed. All firms who develop and use FMs have a responsibility to play their part in addressing risks to consumers, including by increasing consumer awareness and understanding, and the accountability of FMs.

---

[39] Google: The Keyword Blog (2023), A new partnership to promote responsible AI

The market is more likely to develop towards positive outcomes if:

- FM developers and deployers face competitive pressure to improve the reliability and accuracy of their models.

- There is a mechanism to determine the proper allocation of accountability and responsibility.

- Consumers are made aware if content is FM-generated and the risks and limitations associated with FM-generated content, such as whether it is reliable, so they can make informed choices.

- FM developers provide sufficient, understandable and accurate information to businesses, so they understand the relevant characteristics of the models, manage their own risk and prevent harm to consumers.

- FM developers and deployers protect consumers by ensuring that appropriate safeguards are in place to protect people from bad actors using FMs.

We will be vigilant in keeping market developments under review to ensure that consumers are sufficiently protected. This includes by bringing consumer enforcement action where this is appropriate.

Competition and consumer protection law, and the role for regulation

1.85   There will be an important role for regulation as FMs and FM applications develop further. But as with all regulation, it needs to be proportionate and targeted at identified risks. Overly burdensome regulation may make it unnecessarily difficult for competition and innovation to flourish, and at worst may lead to concentration and become a significant barrier to entry in its own right. We have seen evidence indicating that there is currently vibrant competition at some levels of the supply chain (eg in FM development). Where markets are competitive and working well, leaving markets to develop organically may be the best way to achieve positive outcomes for businesses and consumers.

1.86   If we see competition or consumer issues begin to emerge in previously competitive markets, the CMA has a range of existing powers to address these. These include powers to take action against businesses and individuals that take part in cartels or other anti-competitive behaviour, protect consumers from unfair trading practices, and investigate entire markets if we think there are competition

or consumer problems. The CMA will not hesitate to use its powers where appropriate.

1.87 The CMA's ability to protect consumers and promote growth in the UK economy by ensuring free and vigorous competition amongst businesses will be enhanced once the Digital Markets, Competition and Consumers (**DMCC**) Bill comes into force,[40] including through setting targeted conduct requirements on firms found to have strategic market status in respect of a digital activity.

1.88 We look forward to working closely with the government and other regulators, including as part of the Digital Regulation Cooperation Forum (DRCF), to help ensure that the regulation of FMs strikes the right balance between policy objectives, is pro-innovation and delivers the best outcomes for consumers and the UK economy.

1.89 The CMA's focus is on its remit of competition and consumer protection. However, we recognise that these issues cannot be looked at isolation and may interact with other policy objectives. While other policy areas will be the focus of other regulatory authorities, it is important that these different policy areas develop in a joined-up way. Careful thought will therefore be needed when implementing future regulation to take account of a range of policy objectives, including the potential impact of such regulation on effective competition.

Competition and consumer protection principles that can best guide the future development and deployment of FMs.

1.90 To ensure that people, businesses, and the wider economy benefit from the innovation AI can bring, businesses must comply with existing consumer and competition law. Alongside this, to ensure that competition and consumer protection remains an effective driving force as the development and deployment of FMs evolves, we propose the following guiding principles:

---

[40] The DMCC Bill is making progress through Parliament and has now moved to its third reading in the House of Commons.

The diagram organizes principles under three categories on the left, with ACCOUNTABILITY as an overarching theme, and detailed bullet points on the right.

Category	Principle	Details
MODEL DEVELOPMENT	**ACCESS** — Ongoing ready access to key inputs	• Access to data, compute, expertise and capital without undue restrictions. • Continuing effective challenge to early movers from new entrants. • Successful FM developers do not gain an entrenched and disproportionate advantage by being the first to develop a FM, having economies of scale or benefitting from feedback loops. • Powerful partnerships and integrated firms do not reduce others' ability to compete.
MODEL DEVELOPMENT	**DIVERSITY** — Sustained diversity of business models, including both open and closed	• Both open and closed source models push the frontier of new capabilities. • Open-source models help reduce barriers to entry and expansion.
USE OF MODELS IN OTHER MARKETS	**CHOICE** — Sufficient choice for businesses so they can decide how to use FMs	• A range of deployment options, including in-house FM development, partnerships, APIs or plug-ins.
USE OF MODELS IN OTHER MARKETS	**FLEXIBILITY** — Flexibility to switch or use multiple FMs according to need	• Interoperability to support firms mixing and matching or deploying multiple FMs. • Consumers can switch and/or use multiple services easily and are not locked into one provider or ecosystem.
USE OF MODELS IN OTHER MARKETS	**FAIR DEALING** — No anti-competitive conduct, including anti-competitive self-preferencing, tying or bundling	• Confidence that the best products and services will win out. • No anti-competitive conduct, including anti-competitive self-preferencing, tying or bundling, especially from vertical integration. • Competition can counteract any data feedback or first mover effects.
USE OF MODELS BY CONSUMERS	**TRANSPARENCY** — Consumers and businesses are given information about the risks and limitations of FM-generated content so they can make informed choices	• People and businesses are informed of FMs' use and limitations. • Developers give deployers the information to allow them to manage their responsibilities to consumers.

ACCOUNTABILITY — FM developers and deployers are accountable for outputs provided to consumers

1.91 Factors that could undermine these principles include, but are not limited to:

- Mergers or acquisitions which could lead to a substantial lessening of competition in markets for the development or deployment of FMs.

- If firms use their leading positions in key markets to block innovative challengers who develop and use FMs (for example, by misuse of vertical integration).

- Undue restrictions on firms' ability to switch between or use multiple FM providers.

- The development of ecosystems that unduly restrict choice and interoperability.

- If firms with market power in FM development or deployment engage in anti-competitive conduct such as the tying or bundling of products and services.

- If consumers receive false and misleading content from FM services that impacts or is likely to impact their decision-making.

Next steps

1.92   This initial review has been possible as a result of constructive and collaborative inputs from a wide range of people and businesses. We plan to continue the collaborative spirit of our work to date as we take it forward to the next stage. We have proposed this set of principles, but we do not see them as the finished article; instead, we plan to seek views both on report overall and on the principles themselves. This will help ensure that the principles can support the best outcomes for people, businesses and the economy, including through helping firms work to deliver them.

1.93   To that end we are now starting a significant programme of engagement, which will take place in the UK, US and elsewhere over the coming months.

1.94   We plan to speak to a wide range of people to seek views, including:

- Consumer groups and civil society representatives

- Leading FM developers such as Google, Meta, OpenAI, Microsoft, NVIDIA and Anthropic

- Major deployers of FMs

- Innovators, challengers and new entrants

- Academics and other experts

- Government

- Fellow regulators, in the UK including via the Digital Regulators Cooperation Forum, and further afield with our international counterparts.

1.95   We will publish an update on our thinking on the principles, and how they have been received and adopted, in early 2024, also reflecting on further developments in the market.

1.96   We hope that this collaborative and iterative approach will help guide the market to more positive outcomes and realise the maximum potential of this new technology, but we are ready to intervene where necessary.