



FEDERAL TRADE COMMISSION  
PROTECTING AMERICA'S CONSUMERS

Technology Blog

# Generative AI Raises Competition Concerns

By: Staff in the Bureau of Competition & Office of Technology | June 29, 2023 |   

Generative AI has the potential to rapidly transform the way we live, work, and interact. Within just a few months, generative AI chatbots and applications have launched and scaled across industries and reached hundreds of millions of people. AI is increasingly becoming a basic part of daily life.

Generative AI depends on a set of necessary inputs. If a single company or a handful of firms control one or several of these essential inputs, they may be able to leverage their control to dampen or distort competition in generative AI markets. And if generative AI itself becomes an increasingly critical tool, then those who control its essential inputs could wield outsized influence over a significant swath of economic activity.

The FTC's Bureau of Competition, working closely with the Office of Technology, is focused on ensuring open and fair competition, including at key inflection points as technologies develop. Generative AI represents one of these paradigm shifts. Accordingly, it is especially important that firms not engage in unfair methods of competition or other antitrust violations to squash competition and undermine the potential far-reaching benefits of this transformative technology. Unfair methods of competition can distort the rate and direction of innovation. By contrast, open and competitive markets can pave the way for emerging technologies, such as generative AI, to yield their maximum potential benefit.

This blog post identifies a few of the essential technical building blocks of generative AI and discusses competition concerns potentially raised by generative AI.

## What Is Generative AI?

"Generative AI" is a category of AI that empowers machines to generate *new* content rather than simply analyze or manipulate existing data. By using models trained on vast amounts of data, generative AI can generate content—such as text, photos, audio, or video—that is sometimes



indistinguishable from content crafted directly by humans. Large language models (LLMs), which power chatbots and other text-based AI tools, represent one common type of generative AI.

Many generative AI models are developed using a multi-step process: a pre-training step, a fine-tuning step, and potential customization steps. These steps may all be performed by the same company, or each step may be performed by a different company. The pre-training step creates a base model with broad competency in a specific domain, such as language or images. For example, a pre-trained language model might take a partial sentence such as “the family brought their pet goat to the...” and generate potential “autocomplete” suggestions like “park,” “vet,” or “farm.” After pre-training, the model is fine-tuned for a specific application, such as responding to questions or generating images from prompts. In a chatbot interface, a user may ask: “What are some good places to bring your pet goat?” Finally, some types of generative AI can be further customized via methods specific to certain types of models, such as [prompt engineering](#). Prompt engineering is used by many chatbot developers to add more constraints—directions to not respond to inappropriate or harmful questions—or to imitate behaviors.

## The Essential Building Blocks of Generative AI

Generative AI may raise a variety of competition concerns. In particular, control over one or more of the key building blocks that generative AI relies on could affect competition in generative AI markets.

### *Data*

The foundation of any generative AI model is the underlying data. Developing generative AI typically requires exceptionally large datasets, especially in the pre-training step. The data used in this step forms the foundation of the model in the chosen domain, such as language or images.

The volume and quality of data required to pre-train a generative AI model from scratch may impact the ability of new players to enter the market. There are a few reasons why collecting a large and diverse corpus of data can be harder for new market entrants than for incumbents. First, more established companies may benefit from access to data collected from their users over many years—especially if the incumbents also own digital platforms that amass large amounts of data. This data may also be more detailed and robust. Second, established companies are more likely to have developed and honed proprietary data collection tools and technologies for acquiring or scraping data.

This may be particularly true in specialized domains or domains where data is more highly regulated [↑](#)

such as healthcare or finance. Pre-training or fine-tuning a model with deep expertise in these types of areas may require access to large amounts of data that is not widely available and would be difficult for a new player in the market to collect.

Of course, simply having large amounts of data is not unlawful – though the Commission has brought many actions alleging that companies' data collection, retention, or usages policies and practices were unreasonable, unfair, or deceptive. However, even with responsible data collection practices in place, companies' control over data may also create barriers to entry or expansion that prevent fair competition from fully flourishing.

### ***Talent***

Another essential input for generative AI is labor expertise. Developing a generative model requires a significant engineering and research workforce with particular—and relatively rare—skillsets, as well as a deep understanding of machine learning, natural language processing, and computer vision. It can be difficult to find, hire, and retain the talent required to develop generative AI.

Additionally, the speed and velocity at which generative AI is evolving means that models may quickly become outdated or obsolete. The talent companies can acquire and maintain may play a key role in not only the path, but also the rate, of generative AI's evolution.

Firms hoping to compete in the generative AI space need expertise, not only on how to develop generative AI but also on how to deploy the fine-tuned AI products. Companies that can acquire both the engineering experience and professional talent necessary to build and package the final generative AI product or service will be better positioned to gain market share.

Since requisite engineering talent is scarce, powerful companies may be incentivized to lock-in workers and thereby stifle competition from actual or would-be rivals. To ensure a competitive and innovative marketplace, it is critical that talented individuals with innovative ideas be permitted to move freely, and, crucially, not be hindered by non-competes.<sup>[\[1\]](#)</sup>

### ***Computational Resources***

Access to computational resources is a third key input in generative AI markets. Generative AI systems typically require significant computational resources ("compute"). This is especially true at the pre-training step, when creating a new model from scratch. Compute allows generative AI companies to process data, train the model, and deploy the AI system. Compute generally requires  dedicated hardware, such as computers with specialized chips like graphical processing units (GPUs,

that can be expensive to operate and maintain. New entrants typically access compute by turning to cloud computing services, which provide high-performance compute resources on demand. However, cloud services can be expensive and are currently provided by only a handful of firms, raising the risk of anticompetitive practices.<sup>[\[2\]](#)</sup>

The fine-tuning step usually requires substantially less compute than pre-training steps, as fine-tuning techniques usually involve smaller training datasets. For example, the fine-tuning technique LoRa (Low-Rank Adaptation) can enable a developer to fine-tune a model to perform a specific task using consumer grade GPUs. However, fine-tuning must be performed on an existing pre-trained base model, meaning that companies wishing to create a fine-tuned model for a specific application must partner with a company that already owns an established base model, take on the substantial cost of developing their own, or turn to publicly available open-source base models.

This high cost of entry to creating a pre-trained base model may lead to a market where the highest quality pre-trained models are controlled by a small number of incumbents. Several developments currently ongoing in generative AI research may affect this dynamic, such as the proliferation of open-source pre-trained models and [models with a much smaller number of parameters](#), which are cheaper to train. The evolution of the relative quality of public base models compared to privately controlled models will likely be a key factor in determining the impact of computational resources as a barrier to entry in generative AI development.

Additionally, some markets for specialized chips are—or could be, without appropriate competition policies and antitrust enforcement—highly concentrated. Last year [a challenge by the FTC led to Nvidia](#) abandoning its proposed \$40 billion acquisition of Arm. The FTC's complaint alleged that the merger would have stifled competition in multiple processor markets, including chips for cloud service providers. Today, increasing demand for server chips may outpace supply in some instances. There are reports, for example, that the spike in demand for server chips that can train AI has caused a shortage, prompting major cloud-server providers such as AWS, Microsoft, Google, and Oracle to ["limit their availability for customers."](#) And firms in highly concentrated markets are more prone to engage in unfair methods of competitions or other antitrust law violations.

### The Potential Impact of Open-Source

The open-source ecosystem may play an important role in the development of generative AI. For example, last year, following major advances in closed-source image generation models, open-source image generation models became available to the open-source community. What followed was a proliferation of open-source models with similar capabilities. Rapidly, the capabilities of the



open-source image generation models eclipsed those of the proprietary base models that inspired them. The open-source innovation explosion in image generation, coupled with new developments in optimizations, made it possible for nearly anyone to develop, iterate on, and deploy the models using smaller datasets and lower-cost consumer hardware. In this manner, the open-source ecosystem may help open up the playing field once the base models become available to the public, if it can reach parity with the quality of proprietary models. However, open-source AI models are also susceptible to misuse. For instance, while open-source AI image generation tools were released with built-in restrictions on the types of images that could be generated, malicious users removed these protections and utilized the models to create non-consensual intimate images.

Experience has also shown how firms can use “open first, closed later” tactics in ways that undermine long-term competition. Firms that initially use open-source to draw business, establish steady streams of data, and accrue scale advantages can later close off their ecosystem to lock-in customers and lock-out competition.

Products and services using generative AI capabilities may develop in both open-source and proprietary ecosystems. In a proprietary ecosystem, access to the essential building blocks of generative AI—and who controls that access—may play a major role in determining which firms sink or swim.

### **Possible Unfair Methods of Competition**

Incumbents that control key inputs or adjacent markets, including the cloud computing market, may be able to use unfair methods of competition to entrench their current power or use that power to gain control over a new generative AI market.

For example, market leaders could attempt to foreclose competition through bundling and tying. Bundling occurs when a company offers multiple products together as a single package. Tying occurs when a firm conditions the sale of one product on the purchase of a separate product. Incumbents may be able to link together new generative AI applications with existing core products to reduce the value of competitors’ standalone generative AI offerings, potentially distorting competition.

Incumbents that offer a range of products and services as part of an ecosystem may also engage in exclusive dealing or discriminatory behavior, funneling users toward their own generative AI products instead of their competitors’ products. Further, incumbents that offer both compute services and generative AI products—through exclusive cloud partnerships, for instance—might use their power in the compute services sector to stifle competition in generative AI by giving discriminatory treatment 

to themselves and their partners over new entrants. A related scenario exists where an incumbent offers both their own products leveraging generative AI as well as offering APIs allowing other companies to leverage their generative AI capabilities. In such circumstances, there is a risk that incumbent firms will offer their APIs on terms which exist to protect their incumbent position.

Incumbent firms could also use M&A activity in the generative AI space to consolidate market power in the hands of a few players. Large firms already active in generative AI—or that already control a critical input—may try to buy up critical applications and cut off rival access to core products. Market leaders may also try to buy up complementary applications and bundle them together. Additionally, incumbents may be tempted to simply buy up nascent rivals instead of trying to out-compete them by offering better products or services.

### **Network and Platform Effects Can Supercharge Harms from Unfair Conduct**

Firms in generative AI markets could take advantage of network effects to maintain a dominant position or concentrate market power. A first mover could secure a significant advantage over its competitors because its models, by virtue of having interacted with a greater number of users over a longer period, are able to generate more engaging and useful content than rival products. Because positive feedback loops can improve the performance of generative AI models, generative AI products can get better the more people use them. At the same time, this can result in a concentrated market with less possibility for entrants to compete effectively. Absent legal or policy intervention, network effects can supercharge a company's ability and incentive to engage in unfair methods of competition.

Another related effect is platform effects, where companies may become dependent on a particular platform for their generative AI needs. As with network effects, firms could leverage platform effects to consolidate their market power, especially if they take specific steps to lock in customers in an exclusionary or otherwise unlawful way. One specific area where platform effects may play a significant role is cloud services. Cloud providers may exploit generative AI companies' need for compute by trying to lock in customers by, for example, charging exorbitant data egress fees.

### **Looking Ahead**

The FTC is no stranger to dealing with emerging technologies. Generative AI is still evolving rapidly, but it already has the potential to transform many markets. Through vigorous law enforcement, the FTC strives to support a vibrant marketplace where new businesses can compete, researchers are free to move to the jobs where they can best advance the state of the technology, and entrepreneur



can continue to innovate. As competition issues surrounding generative AI continue to develop, the Bureau of Competition, working closely with the Office of Technology, will use our full range of tools to identify and address unfair methods of competition.

----

Thank you to the contributors of this post: Elena Ponte, John Newman, Dan Principato, David B. Schwartz, Jake Walter-Warner, David Koh, Alex Gaynor, Nick Jones, Stephanie Nguyen, Ben Swartz, Varoon Mathur, Josephine Liu, Daniel Zhao, Dan Salsburg, Sam Levine.

---

[\[1\]](#) The FTC recently issued a notice of proposed rulemaking that could potentially help to address such concerns. Non-Compete Clause Rule, 88 Fed. Reg. 3482 (proposed Jan. 19, 2023), <https://www.federalregister.gov/documents/2023/01/19/2023-00414/non-compete-clause-rule>.

[\[2\]](#) The FTC recently conducted an inquiry into the business practices of Cloud Computing Providers, which may also further address the risk of anticompetitive practices. See Solicitation for Public Comments on the Business Practices of Cloud Computing Providers, [FTC-2023-0028](#).

**Tags:** [Technology](#) | [Artificial Intelligence](#) | [Office of Technology](#)

#### More from the Technology Blog

---

Technology Blog

## The FTC Is on the Front Lines of Tech Innovation & Regulation

Stephanie T. Nguyen, Chief Technologist | January 17, 2025

Technology Blog

## Behind the FTC's 6(b) Report on Large AI Partnerships & Investments

Office of Technology Staff | January 17, 2025

Technology Blog

## Surveillance Pricing Update & The Work Ahead

Stephanie T. Nguyen & Samuel A.A. Levine | January 17, 2025



Technology Blog

## Lenses of security: Preventing and mitigating digital security risks through data management, software development, and product design for humans

Staff in the Office of Technology & the Division of Privacy and Identity Protection | December 13, 2024

Get Business Blog updates

Subscribe

