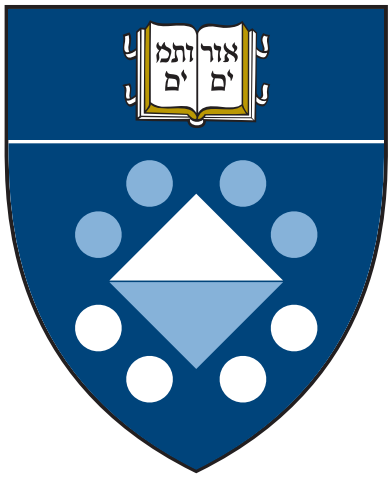

Yale University

Thurman Arnold Project

2024 Student Projects

Paper 8

Generative AI and Competitive Threats

May 2024

Kart Kandula
Brandon Nye
Ruth Zheng

Generative artificial intelligence (AI) models have emerged at the forefront of recent policy debates. Generative AI has the potential to multiply productivity, but like many technological advancements in the past few decades, it also risks putting enormous power in the hands of a few dominant firms. Our paper surveys several threats that generative AI poses to competition, consumers, and society, and recommends potential policy solutions. Part I discusses the existence of economies of scale and vertical integration along the AI tech stack, and the resulting tendency for concentration and exclusion. Part II explores generative technologies' heightened risks of bias and self-preferencing. Part III focuses on the risk that AI poses to online news publications and other content creators.

I. Economies of Scale, Vertical Integration, and Tendencies to Exclude

The generative AI market, much like the search engine market, will tend toward monopoly and oligopoly. There exist economies of scale in the development of generative AI models, and so far, there has been vertical integration along the tech stack in generative AI. Together, these two phenomena create the incentive and ability for generative AI firms to engage in exclusionary conduct. This will not only undermine consumer choice but also stymie innovation.

A. Generative AI experiences economies of scale

There are economies of scale in generative AI. The pre-training phase, where foundational models are molded using vast troves of data typically collected from the internet, has received considerable attention over data appropriation and intellectual property issues.[1] During this pre-training phase, large language models (LLMs) develop a base of general language understanding and a degree of factual knowledge. Without further refinement, however, these models are inherently limited, only able to predict subsequent words in a text sequence. Less focus has been placed on the refinement stage. While pre-training occurs in the development stage, refinement through a process called *fine tuning* occurs downstream during the deployment phase.[2]

Fine tuning refines the model's weights, enabling it to engage in conversations with users and aligning it with desirable human conduct. The fine-tuning phase utilizes a labeled dataset, which includes examples of both desirable and undesirable outputs, to instruct the model explicitly on target outcomes and behaviors. In layman's terms, as users tell LLMs like AI chatbots whether the output is helpful or unhelpful, the model continues to learn and improve its responses.[3] The more refined and helpful the model's responses, the more consumers will engage with the chatbot. There is thus a positive feedback loop. As models are increasingly deployed and interact

---

[1] *See, e.g.*, Michael M. Grynbaum & Ryan Mac, *The Times Sues OpenAI and Microsoft Over A.I. Use of Copyrighted Work*, N.Y. TIMES (Dec. 27, 2023), https://www.nytimes.com/2023/12/27/business/media/new-york-times-open-ai-microsoft-lawsuit.html; Kyle Chayka, *Is A.I. Art Stealing from Artists?*, THE NEW YORKER (Feb. 10, 2023), https://www.newyorker.com/culture/infinite-scroll/is-ai-art-stealing-from-artists.

[2] Information Technology Industry Council (ITI), *Understanding Foundation Models & The AI Value Chain: ITI's Comprehensive Policy Guide* (Aug. 2023), https://www.itic.org/documents/artificial-intelligence/ITI_AIPolicyPrinciples_080323.pdf.

[3] C. Scott Hemphill, *Disruptive Incumbents: Platform Competition in an Age of Machine Learning*, 119 COLUM. L. REV. 1973, 1978 (2019) (noting that "the importance of scale is heightened by the high variability of user data" in LLMs).

with diverse user inputs, the firm collects new data on the desirability of outputs. This new data, once integrated back into the fine-tuning process, enhances the model's understanding and responsiveness, creating a feedback loop that continuously improves performance. Just as Google's amassed click-and-query data gives it a major competitive edge over its smaller search rivals because the data is used to improve the relevance of the search results, [4] the generative AI model with the most user responses will experience better refinement, giving it a competitive advantage. These "significant data network effects… giv[e] rise to significant first-mover advantages."[5] Thus, like search, the generative AI market enjoys economies of scale and tends toward monopoly.[6]

B. Significant concentration along the AI tech stack

The perception of vigorous competition in generative AI, especially at the foundation model and application levels,[7] belie high levels of concentration along the full technology stack. Four general layers make up this stack: 1) hardware, which is the infrastructure undergirding computation; 2) the cloud, which sustains the computation required to train and run the models; 3) the models themselves; and 4) user-facing applications.[8]

First, the hardware layer experiences an extreme level of concentration. Generative AI is undergirded by special chips, Graphics Processing Units (GPUs), that are specifically designed to perform intensive, technical calculations at speed. Analysts have estimated that over 90% of these chips are supplied by Nvidia.[9] The past few years have witnessed steep price increases and supply constraints, giving rise to bottlenecks and hoarding.[10] Second, the fixed costs for data and equipment are even higher for generative AI than traditional software development, while marginal costs remain comparatively low, suggesting natural monopoly conditions.[11] Training

---

[4] *See generally*, Lina Khan, *The Separation of Platforms and Commerce*, 119 COLUM. L. REV. 973 (2019) (highlighting the self-reinforcing advantages of data which may give incumbents a significant enough lead such that new potential competitors are rendered unable to enter).

[5] Tejas N. Narechania & Ganesh Sitaraman, *An Antimonopoly Approach to Governing Artificial Intelligence*, YALE L. & POLICY REV. (forthcoming 2024) (manuscript at 18).

[6] It is not necessarily inevitable that the generative AI market has this tendency. In the search engine context, scholars have proposed imposing interoperability requirements, including data sharing, in order to limit the extent to which search engines may enjoy a natural monopoly. *See, e.g.*, Herbert Hovenkamp, *Antitrust Interoperability Remedies*, 123 COLUM. L. REV. 1 (2023). A similar remedy would reduce the economies of scale in the generative AI market.

[7] *See* Shana Lynch, *AI Index: State of AI in 13 Charts*, Stanford HAI (Apr. 15, 2024), https://hai.stanford.edu/news/ai-index-state-ai-13-charts (finding organizations released 149 foundation models in 2023).

[8] Tejas N. Narechania & Ganesh Sitaraman, *An Antimonopoly Approach to Governing Artificial Intelligence* (2023) (forthcoming); *see also* Competition & Markets Authority (CMA), AI Foundation Models Update Paper (Apr. 11, 2024) https://assets.publishing.service.gov.uk/media/661941a6c1d297c6ad1dfeed/Update_Paper__1_.pdf.

[9] Philipp Wegner, *The Leading Generative AI Companies*, IOT ANALYTICS (Dec. 14, 2023), https://iotanalytics.com/leading-generative-ai-companies/.

[10] *Id.*; *see also* Reuters, *EU Examines Nvidia-Dominated AI Chip Market's Alleged Abuses, Bloomberg Reports*, REUTERS (SEP. 30, 2023, 5:20 PM EDT), https://www.reuters.com/technology/eu-starts-early-stage-probe-into-nvidia-dominated-ai-chip-market-abuses-2023-09-29/; Aaron Mok, *Meta Is Sitting on a Stockpile of Tech's Hottest Commodity*, BU. INSIDER (Jan. 19, 2024, 2:02 PM EST), https://www.businessinsider.com/meta-zuckerberg-stockpile-tech-hottest-commodity-nvidia-chips-semiconductors-2024-1 ("Meta is hoarding chips as it moves toward being an AI-first company.").

[11] Tejas N. Narechania, Machine Learning as Natural Monopoly, 107 IOWA LAW REVIEW 1543, 1576-78 (2022).

costs have skyrocketed in the past years; whereas the original Transformer model cost about $900 to train, GPT-4 cost an estimated $78M and Gemini Ultra $191M.[12] Third, the cloud itself is mostly controlled by three dominant players: Amazon (AWS), Microsoft (Azure), and Google (GCP). Tejas Narechania and Ganesh Sitaraman have expressed concern that a small number of already powerful technology companies are active at every layer of the stack.[13] Microsoft is investing in developing its own AI chips, runs the Azure native public cloud platform, has a massive stake in OpenAI, and runs a host of end-user applications.[14] Google is developing its own chips, runs GCP, has developed its own foundation models, and runs end user applications like Google Bard.[15] Amazon has announced a strategic partnership with Nvidia,[16] runs AWS, has recently bought a major stake in Anthropic (a rival to OpenAI),[17] and runs applications.

C. Competitive harms

a. Exclusionary unilateral conduct

Last month, the UK Competition & Markets Authority identified three risks to competition posed by generative AI, which we elaborate on.[18] First is the risk that firms controlling critical inputs to development along the tech stack will engage in exclusionary conduct to block entry and stifle nascent competition.[19] This would be similar to Microsoft's agreements with PC manufacturers to retain its browser monopoly[20] and the recent charges against Google for its agreements with mobile manufacturers to maintain its search monopoly.[21] A generative AI firm could cut off access to chips, data, or cloud instances or otherwise condition access on making its foundation model the default. Like Google's Android strategy, a firm could condition usage of its open-source models on anti-forking agreements.[22] This would raise rivals' already high costs to develop competing foundation models, thus shielding the top players from competition.

b. Self-preferencing content or applications

Second, on the flip side, incumbents might exploit their downstream access to consumers and businesses to self-preference their related-market offerings – content, applications, products, etc. There is robust evidence that Google does this with search, ranking its own offerings above competitors in key verticals like travel and requiring competitors to pay their way back to the top

---

[12] Lynch, *supra* note 7.
[13] Narechania & Sitaraman, *supra* note 5.
[14] *Id.* at 21.
[15] *Id.* at 20.
[16] Steve McDowell, *AWS and Nvidia Expand AI Relationship*, FORBES (Mar. 20, 2024, 10:44 PM EDT) https://www.forbes.com/sites/stevemcdowell/2024/03/20/awss-surprisingly-deep-nvidia-relationship/.
[17] Alex Heath, *Amazon Scrambles for its Place in the AI Race*, THE VERGE (Mar. 29, 2024, 7:45 PM EDT), https://www.theverge.com/2024/3/29/24116056/amazon-ai-race-anthropic-olympus-claude.
[18] CMA, CMA AI Strategic Update (Apr. 29, 2024), https://www.gov.uk/government/publications/cma-ai-strategic-update/cma-ai-strategic-update#the-cmas-understanding-of-the-risks-posed-by-ai.
[19] *Id.*
[20] *See generally*, United States v. Microsoft Corp., 253 F.3d 34, 51, 82 (D.C. Cir. 2001) (en banc) (per curiam).
[21] Complaint, Colorado v. Google, No. 1:30-cv-03715 (D.D.C. Dec. 17, 2020).
[22] *Id.*

by purchasing ads.[23] Generative AI creates an even stronger incentive and ability to self-preference. Unlike traditional search, generative AI applications will typically output a summary of results in response to a query rather than a list, often without linking the original sources. Coupled with the fact that self-preferencing may be baked into various layers of the algorithm itself, this makes it difficult for self-preferencing to be detected. Self-preferencing and bias will be discussed at a deeper level in Part II. And issues of appropriation will be discussed in Part III.

c. Exclusionary contracts and partnerships

Finally, and related to the first two risks, incumbents and key players may forge partnerships to maintain their market power.[24] For one, Apple and Google have already been in discussion to make Google's AI model, Gemini, the default on iPhones.[25] This would serve to extend Apple and Google's longstanding search partnership. Amazon's partnership with Nvidia might be another example. Given the scarcity of Nvidia's flagship GPUs, such agreements could be understood as Nvidia doling out preferential treatment to customers in exchange for strategic alignment.[26] There is significant incentive for these strategic partnerships to be wielded to exclude rivals from key inputs or access to downstream markets, ultimately to the detriment of long run innovation and healthy competition.

D. Policy proposals

Many scholars and regulatory bodies have suggested that ex post antitrust enforcement against generative AI firms will need to be coupled with ex ante regulation. [27] As such, competition agencies may consider advocating for a digital regulator. While AI may be the right hook, such a regulator should aim to cover digital markets broadly to bring U.S. competition

---

[23] *Id.* This conduct is particularly pronounced in travel, shopping, and local. *See* Khan, *supra* note 4 ("In Europe and India, competition authorities have found that Google ranks its own services higher than those offered by rivals, a 'search bias' that means anyone competing with Google properties may effectively disappear from Google search results.").

[24] CMA *supra* note 18; *see also* Barry Lynn, Max von Thun, & Karina Montoya, *AI in the Public Interest: Confronting the Monopoly Threat*, OPEN MKTS. INST. (November 2023), https://static1.squarespace.com/static/5e449c8c3ef68d752f3e70dc/t/6554461d58cc944a2d95bc6e/1700021790820/OMI+AI+Report+WEB.pdf.

[25] Tripp Mickle, Nico Grant, & Brian X. Chen, *Apple and Google Are Discussing a Deal to Bring Generatve A.I. to iPhones*, N. Y. TIMES (Mar. 19, 2024), https://www.nytimes.com/2024/03/19/technology/apple-google-ai-iphone.html.

[26] McDowell, *supra* note 16 ("Nvidia is not a neutral party in this, favoring partners willing to embrace its platform-focused strategy.").

[27] The CMA report from April highlights six key principles for enforcement and regulation: open access to key inputs; diversity of models; choice between models, including easy switching; fair dealing, meaning no self-preferencing, tying, or exclusionary contracts; transparency not just for consumers and businesses but also deployers; and accountability. *See* Narechania & Sitaraman, *supra* note 5.

policy in line with counterparts in the E.U.,[28] U.K.,[29] and likely soon Japan[30] and Australia.[31] Like the Digital Markets Act in the EU and the Digital Markets, Competition and Consumers Bill in the UK, authorities could seek to segment the generative AI market, targeting the most threatening firms based on size or market power.

More specifically, to guard against exclusionary conduct, competition authorities can impose ex ante requirements for open access, fair dealing, and interoperability through guidance or rulemaking.[32] More forceful measures to foster openness and contestability might include requirements to share data. Authorities can also police cancellation mechanisms at each level of the AI tech stack to ensure that switching between applications, foundation models, cloud providers, and hardware is easy and low cost.[33]

To address self-preferencing, competition authorities could impose transparency requirements forcing generative AI applications to provide disclaimers, link sources, or disclose decision parameters. Regulators could also require AI applications to offer users the opportunity to toggle relative input weights. This might also make sense as part of a guidance document or general AI rule. A digital regulator may also need to build out the independent authority to audit algorithms and training data and screen for systematic forms of bias. In addition, it may be sensible to impose responsibility onto generative AI companies themselves to publish reports that document affirmative action taken to mitigate the perceived risk of self-preferencing. This will be discussed further in Part II.

To guard against anti-competitive collusion, competition agencies should open investigations into contracts in the space and bring Section 2 enforcement actions against firms where contractual terms may serve to block entry. The next section will deep dive into self-preferencing and bias.

II. Self-preferencing and algorithmic bias

A. Heightened risks of undetected self-preferencing

Due to high levels of vertical integration and the structure of many generative AI applications, firms may be able to engage in even more egregious forms of undetected self-preferencing than in traditional search. Unlike traditional search engines, which allow users some degree of choice

---

[28] *See* Council Regulation 2022/1925, 2022 O.J. (L 265) 1 [hereinafter Digital Markets Act].
[29] *See* COMPETITION & MKTS. AUTH., DIGITAL MARKETS UNIT (Apr. 7, 2021), https://www.gov.uk/government/collections/digital-markets-unit.
Press Release: New Watchdog to Boost Online Competition Launches, U.K. Digital Markets Unit.
[30] Michael Mandel, *What the DMA Experiment Means for Japan*, PROGRESSIVE POLICY INST. (Mar. 26, 2024), https://www.progressivepolicy.org/blogs/what-the-dma-experiment-means-for-japan/.
[31] Reuters, *Australian Regulator Calls for New Competition Laws for Digital Platforms*, REUTERS (Nov. 27, 2023 4:07 PM EST), https://www.reuters.com/technology/australian-regulator-calls-new-competition-laws-digital-platforms-2023-11-27/.
[32] *See, e.g.*, Digital Markets Act, *supra* note 28.
[33] See, for example, the FTC's investigation of Amazon for its Prime cancellation protocol. *FTC Takes Action Against Amazon for Enrolling Consumers in Amazon Prime Without Consent and Sabotaging Their Attempts to Cancel* (Jun. 21, 2023), https://www.ftc.gov/news-events/news/press-releases/2023/06/ftc-takes-action-against-amazon-enrolling-consumers-amazon-prime-without-consent-sabotaging-their.

among a list of results—even if, in practice, users' attention rarely extends past the top-ranked results—generative technologies streamline the process by providing a single synthesized answer. Rather than forcing a user to sift through a list of ranked search results, AI Chatbots output a summary of the results.[34] Further, most models offer the benefit of conforming their query responses to what they predict a user's intent to be. One of GPT-4's stated improvements is that it is being "much better at following user intentions."[35] Users are wont to accept the responses and move on. Indeed, this convenience is precisely the value proposition that draws users to generative AI applications.

However, this feature that enables consumers to save time by offloading a degree of information processing and decision-making agency to an AI tool should also be a source of potential concern for regulators. Users risk being systematically directed to draw conclusions or make purchasing decisions that are not truly reflective of their individual preferences. This risk is exacerbated by the opacity of not only the decisional criteria being used but also the algorithm employed to predict a user's preferences.[36] This shift from ranked search results to summary output transfers control from users to the model itself. As a result, companies are able to manipulate user experiences by preferencing certain sources or information that align with their own interests or commercial objectives, rather than prioritizing the user's needs or the objective relevance of information.

And it is well documented that generative AI firms have this power. Bias can be systematically built into a foundation model during the pre-training phase.[37] Developers argue that any bias arises not from the algorithm itself but from the underlying data scraped from the internet.[38] Thus it becomes difficult to assign responsibility for biased AI output.

In sum, this greater delegation of decision-making power coupled with the near total opacity of models' actual mechanics increases the risk that self-preferencing may occur without detection. Noga Shchory and Michal Gal discuss this phenomenon with voice shoppers (e.g., Amazon's Alexa).[39] By shifting the decision point from which item to purchase to the binary choice of whether to purchase the recommended item, voice shoppers create consumer choice gaps.[40] This is further exacerbated by the lack of information about the voice assistant's decisional parameters and relative weights.[41] Voice shoppers are thus able to manipulate user experiences by preferencing certain sources or information that align with their own interests or commercial objectives, rather than prioritizing the user's needs or the objective relevance of information.

---

[34] Some chatbots offer inline citations (e.g., Perplexity), but most do not.
[35] Ryan Lowe & Jan Leike, *Aligning Language Models to Follow Instructions*, OpenAI (Jan. 27, 2022), https://openai.com/research/instruction-following.
[36] *See* Long Ouyang et al., *Training Language Models to Follow Instructions with Human Feedback*, arXiv:2203.02155 (Mar. 4, 2022), https://arxiv.org/abs/2203.02155.
[37] *See, e.g.*, Margot E. Kaminski & Jennifer M. Urban, *The Right to Contest AI*, 121 Colum. L. Rev. 1957 (2021); Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 Cal. L. Rev. 671 (2016).
[38] *Id.*
[39] Noga B. Shchory & Michal S. Gal, *Voice Shoppers: From Information Gaps to Choice Gaps in Consumer Markets*, 88 Brooklyn L. Rev. 111 (2022).
[40] *Id.*
[41] *Id.*

Like voice shoppers, AI chatbots are structurally set up to exploit consumers' bounded rationality, allowing operators to serve their own preferences—which may include suggesting the operator's own products or those listed by vendors who pay the operator a fee—at the expense of consumer choice. This systematic bias can be entrenched in subtle ways making it difficult to assign accountability.

B. The grounding phase as an avenue for built-in bias

For instance, although AI providers have limited control over the outputs that derive solely from internal model weights, *the grounding phase* of AI development provides an avenue for fine-grained control over what content is sourced for summarization or synthesis in the AI application's output. Grounding is a process through which deployed AI models access and incorporate external, real-time information from the internet to enhance the accuracy and relevance of their responses.[42] An LLM without grounding functions is like a person who must rely solely on their existing knowledge to answer questions. In contrast, an LLM equipped with grounding capabilities is akin to a person who is permitted to consult a range of external sources before providing a response. Grounding enables the operators of these models to dictate the sources that the model may access, offering a means to manage the quality and scope of the information used. In the grounding phase, an AI application can be developed to preferentially access and disseminate information from online sources owned by its developer or from entities with which the developer has business relationships. This control can ensure that the information conforms to certain standards or biases, significantly influencing the model's outputs.

C. Competitive harms in related markets

---

[42] Grounding is especially critical as chatbots / LLMs aspire to function as the new generation of search engines. Grounding is crucial for maintaining the relevance and accuracy of the information provided by LLMs in a rapidly changing world. Unlike the static knowledge embedded in their internal network weights, grounding allows these models to access up-to-date information, reflecting the latest developments and current events. This dynamism is essential for models tasked with providing information that changes frequently, such as news updates, scientific advancements, or market trends. Furthermore, grounding serves as a critical check against the tendency of LLMs to fabricate details or 'hallucinate.' By anchoring responses to real-time, verifiable data, these models are less likely to produce misleading or factually incorrect information.

These advancements represent a pivotal change in the paradigm of information acquisition and dissemination. Traditionally, search engines operated as gateways to the open internet, indexing the web's content and directing users to external sources. The trajectory of search engine evolution, however, has been bending towards direct information delivery for some time. For instance, search engines like Google have long been refining their algorithms and presentation formats to provide immediate answers to queries within search results pages themselves, through features like knowledge panels, featured snippets, and direct answers.

In this vein, LLMs grounded in real-time internet data represent the next leap in search engine evolution. By synthesizing and presenting information directly in a conversational format, LLMs streamline the search process considerably. This approach can not only make information retrieval more efficient but also transform it into a more user-friendly, interactive experience. Users can engage in a dialogue with a chatbot, probing deeper with follow-up questions, and receiving information that feels personalized and immediate. The practical realization of this innovation is already manifesting in developments such as with OpenAI's ChatGPT and Microsoft's Copilot. ChatGPT, for instance, has evolved to include capabilities for web browsing, enabling it to pull in real-time data for more accurate and current responses. Similarly, by integrating the Bing search engine with its Copilot chatbot, Microsoft is blurring the lines between traditional search engines and AI-driven conversational agents.

This raises questions about what sources the foundation model uses and whether the corporation itself owns or has rights to complementary assets. Here, we return once again to the fact that the major foundation models are owned by gatekeeper digital platforms who manage large-scale digital business, which include the deployment of these AI models to end users. An AI model developed by a company with a diversified portfolio of digital services could disproportionately draw upon these services during the grounding process. This could manifest in the AI model frequently citing or referencing the company's news platform, financial tools, or other proprietary services in its responses. Likewise, the grounding process could prioritize external entities with which the AI developer has commercial relationships (business partnerships, advertising agreements, etc.), leading to a scenario where the AI model preferentially sources information from these affiliated sources. This scenario is not merely theoretical: the integration of commercial considerations into search algorithms and content curation is a well-established practice in the digital domain.[43] The complexity and opacity of AI algorithms can exacerbate this issue, as the rationale behind source selection and prioritization is often not transparent to the end user. And the requirement that end users trace the original sources would defeat the convenience factor that drew them to the application in the first place.

This dynamic empowers generative AI technologies to act as gatekeepers in our digital economy, effectively picking winners and losers among content creators, businesses, and information sources. By preferentially sourcing and amplifying content from affiliated or commercially advantageous sources, these AI models can single-handedly manipulate related markets. Much like the DOJ has argued in its case against Google search, self-preferencing in generative AI applications threatens to raise competitors' costs in these related markets. AI companies could require competitor publishers to pay advertising fees or meet certain conditions if they want to appear in the output more frequently, just as Amazon has done with its Buy Box and Google has done by charging specialized vertical search providers to appear above the organic links.[44] Chatbot output could also feature the AI company's own related-market service offerings rather than those of competitors, just as the FTC has accused Amazon of doing in prioritizing its first-party products in its retail search results.[45] This conduct would reduce competition in publishing and related-market services, which raises prices and limits choices for consumers.

D. Potential solutions

Preferential contracts with third parties should be banned, as they are under the E.U.'s Digital Markets Act.[46] Due to the difficulty of detection, advocacy for a digital regulator may be especially necessary to adequately address the challenges of self-preferencing and bias. In addition to investigation of contracts with commercial partners for evidence of preferential treatment, an

---

[43] *See, e.g.,* Fiona Scott Morton, *Improving the Contestability of E-Commerce in Two Jurisdictions: The Amazon Case*, BRUEGAL (Dec. 4, 2023), https://www.bruegel.org/policy-brief/improving-contestability-e-commerce-two-jurisdictions-amazon-case; Complaint, Colorado et al. v. Google, *supra* note 21.
[44] *Id.*
[45] Complaint, FTC v. Amazon.com, No. 23-01495 (W.D. Wash. Sep. 26, 2023).
[46] Scott Morton, *supra* note 43.

expert digital regulator could conduct regular audits of an algorithm's grounding process in particular to look for evidence of anticompetitive bias.

Further, like the CMA report suggests, it may also be advisable to force digital gatekeepers themselves to share responsibility in ensuring fair conduct, for instance requiring them to proactively publish reports documenting affirmative steps taken to ameliorate the perception of self-preferencing and explaining to government officials and the public how their algorithms work. Authorities can impose transparency requirements forcing generative AI applications to provide disclaimers, link sources, and disclose decision parameters. These companies might also be required to give users the option to manually adjust the input sources that go into the grounding process. For instance, a news-focused AI application could allow users the option to choose for themselves which underlying publications they would like the interface to opt in or out of drawing information from.

If continuous enforcement of a prohibition against self-preferencing would be too difficult in practice, competition agencies may also consider advocating for the imposition of separations regimes—like in railroads, banking, and telecommunications.[47] Such a regime would prevent emerging generative AI firms from entering certain related markets where there is a substantial risk of distortion. It would also force vertically integrated gatekeepers to divest their generative AI business lines or to operate them under a separate entity with significant firewalls between it and the main firm. In addition to heightened scrutiny of new mergers in the AI realm and related markets by competition agencies, a digital regulator could oversee and enforce the breaking up big vertically integrated AI firms.

III. Threats to the Market for Publishing

A. Data appropriation threatens online new publishers and other content creation markets

Generative AI will undoubtedly have a sweeping impact on society. One area in which it poses a particularly pressing threat is the market for online news publication and content creation. Generative AI applications engage in data appropriation in both the pre-training and deployment phases. Firms deploying AI will have the ability and incentive to keep user traffic on their own pages which risks obfuscating original publications and deterring original content creation.

The revenue model for many generative AI applications will likely mirror that in online search.[48] Applications will allow users to engage without charge while selling advertisements. As such, these AI applications will have a strong incentive to keep traffic on their own pages instead of directing them elsewhere in order to collect more ad revenue. In practice, generative AI output typically does not direct users to the websites whose content was used to inform the output.[49] This creates a particularly acute issue in news, in which publications are primarily differentiated based on the quality of reporting which is deeply related to informational fact finding. Once those facts are found, however, it will make little difference to most users where they are getting the information. Even where source material is actually linked by a generative AI application in

[47] *See* Khan, *supra* note 4.
[48] Complaint, Colorado v. Google, *supra* note 21.
[49] *See supra* note 34.

responding to an informational news-related query, users will have even less incentive to actually click on the links and read further. And that is just at the deployment level. The application will certainly not be able to direct users to the underlying training data collected from the internet even if it wanted to.

By appropriating original content to create an application that then competes with the same content for web traffic, generative AI applications risk decimating online news publishers' (and other digital content creators') revenue models which are heavily dependent on selling advertising space. Original publications' websites will not receive greater web traffic and therefore ad, subscription, and affiliate revenue when users view AI output derived from their reporting. At scale, this will severely undermine competition in the market for online news publishing and disincentivize the creation of new content, with severe social consequences.

B. Copyright law is an inadequate protection

Advocates have pursued copyright claims against generative AI applications. The *New York Times*, Getty Images, and a number of authors have each sued AI companies for training their AI applications on their original content in violation of their copyright. The plaintiffs allege that the AI companies are using the plaintiffs' work to build a competitor product, free riding off of their enormous upfront investment in creating the original content—for instance reporting breaking news in high-risk zones—in violation of the Copyright Act.[50]

Case law indicates that these claims are controlled by federal copyright law. In a case presenting a similar fact pattern, Genius sued Google for copying and posting Genius's content in Google's search results panel.[51] Posting Genius's content in this way would allow Google to capitalize on the content while diverting traffic away from Genius's website. The Second Circuit, however, affirmed the district court's ruling that the unfair competition claims by Genius are equivalent to a federal copyright claim.[52] Genius thus has to bring its case against Google under copyright law.

The problem, however, is that copyright law is not suitable for handling many of these claims.[53] Generative AI algorithms operate as probabilistic models in that they use an immense amount of text data to predict probabilistically the most appropriate next word in the output. The algorithms do not republish the text from online publications word-for-word. As a result, it would be hard for online publishers like the *New York Times* to claim that AI companies are "copying" the "constituent elements" of the original work, as is required for proving copyright infringement.[54] Likewise, there is likely a strong fair use defense available.[55]

---

[50] Complaint, N.Y. Times Co. v. Microsoft Corp., No. 23-CV-11195 (S.D.N.Y. Dec. 27, 2023).
[51] ML Genius Holdings LLC v. Google LLC, No. 20–3113, 2022 WL 710744 (2d Cir. Mar. 10, 2022).
[52] *Id.* at *4.
[53] In the *Genius* case, for example, Genius did not own a copyright to the song lyrics on its website. *Id.* at *4 n.3. Thus, it could not bring a copyright suit against Google for republishing those lyrics, even though Genius alleged that Google copied them from Genius's website. *Id.*
[54] Feist Publications, Inc. v. Rural Tel. Serv. Co., Inc., 499 U.S. 340, 361 (1991).
[55] *See, e.g.*, Author's Guild v. Google, 804 F.3d 202 (SDNY 2013); Michael D. Murray, *Generative AI Art: Copyright Infringement and Fair Use*, 26 SMU Sci. & Tech. L. Rev. 259 (2023).

C. Legal alternatives to copyright under unfair competition

It may, therefore, be necessary for competition authorities and courts to carve out a new justiciable standard for AI's infringement: one based not on "copying" protected work but on principles of unfair competition. One analog is the "hot news" doctrine, created by the United States Supreme Court in *Intl. News Serv. v. Associated Press*.[56] In that case, the Court held that it constitutes "unfair competition" for one newspaper to misappropriate and publish the news gathered by a competitor newspaper. While the "ideas" or "facts" of the news are not copyrightable—and thus a newspaper does not have a property right enforceable against the public at large—it is "contrary to good conscience" to allow a competitor to profit off of the fruits of one newspaper's costly efforts "to the disadvantage" of that newspaper.[57]

While lower courts have severely limited the "hot news" doctrine as a matter of tort law,[58] competition enforcers should revive misappropriation doctrines in enforcing federal antitrust law. The Supreme Court has clearly articulated that it is unfair competition to allow a publisher to use the efforts of a competitor to profit at the expense of that competitor. This conduct is especially concerning when it disincentivizes publishing in the first place. Indeed, the FTC has recently emphasized its intention to use its authority to prevent unfair business practices related to copyright in AI markets.[59]

AI is the right test case for this new legal standard. Generative AI has the potential to steal the market share of key newsgathering organizations, authors, and content creators and to do so using those parties' original labor. While misappropriation claims of this sort are challenging because it is hard to draw clear lines, competition authorities can make a strong case that generative AI applications are *sui generis* because of their internalized network effects. There may not be a major theoretical difference between ChatGPT and a blogger who reads published articles on current events, synthesizes them, and posts a summary online. However, functionally, generative AI, like Google search, internalizes network effects because the more consumers use the application, the more intelligent the algorithm becomes. The more intelligent the algorithm becomes, the more consumers are drawn to the product. As with Google search, this market dynamic mimics a natural monopoly. A natural monopoly in publishing would impose significant societal costs. Not only are there democracy concerns with centralizing the source of information

---

[56] 248 U.S. 215 (1918).

[57] *Id.* at 240.

[58] Specifically, the Second Circuit held that the "hot news" doctrine was limited to time-sensitive information. National Basketball Ass'n v. Motorola, Inc., 105 F.3d 841, 845 (2d Cir. 1997). The Second Circuit also held that similar misappropriation doctrines in state law were preempted by the 1976 Copyright Act. *Id.* at 851. Additionally, Judge Posner for the Seventh Circuit held that the "hot news" doctrine was based on an "abandoned" view that federal courts could "formulate common law principles in suits arising under state law." Confold Pacific, Inc. v. Polaris Indus., Inc., 433 F.3d 952, 960 (7th Cir. 2006).

[59] *See* Comment of the United States Federal Trade Commission to the United States Copyright Office, *Artificial Intelligence and Copyright*, Docket No. 2023-6 ("Conduct that may violate the copyright laws . . . may also constitute an unfair method of competition or an unfair or deceptive practice, especially when the copyright violation deceives consumers, exploits a creator's reputation or diminishes the value of her existing or future works, reveals private information, or otherwise causes substantial injury to consumers.").

under one private entity, but also crucial sectors of the economy and the arts—newsgathering,[60] authorship, and content creation of all kinds—will be "substantially threatened."[61]

As a matter of policy, a digital regulator may be necessary to oversee the market for online news publication in addition to the use of legal tools. Such a regulator could work to ensure that generative AI applications are abiding by codes of fair conduct including linking original source material in their outputs, forcing applications to take users directly to the original source if enough of its response would be based on that content, and potentially overseeing some sort of profit-sharing mechanism. Ultimately, where traditional legal and antitrust tools fall short, authorities can advocate for the creation of such a regulator with the ability to create a properly working incentive scheme to continue rewarding the socially vital production of quality reporting.

---

[60] *See* Comment of News Corporation to the United States Copyright Office, *Artificial Intelligence and Copyright*, Docket No. 2023-6.

[61] *See National Basketball Ass'n*, 105 F.3d at 845.