What is the Relevant Product Market in AI?

Lazar Radic[1]

Kristian Stout[2]

Introduction

Artificial Intelligence ("AI") has taken the world by storm, and competition law[3] is no exception.[4] Policymakers, academics, and commentators are struggling to make sense of how to apply competition law principles to burgeoning AI markets.[5] The question is spurred by an impending sense that inaction is likely to lead to monopolistic outcomes that will later be impossible to revert.[6] What is feared is that AI will become dominated by a few large technology companies and, more spuriously, that these will be the same companies that already control vast swathes of the so-called digital sphere.[7] In other words: it will make big tech even bigger.

---

[1] Senior Scholar for Competition Policy, International Center for Law & Economics and Assistant Professor of Law, IE Law School.

[2] Director of Innovation Policy, International Center for Law & Economics.

[3] The terms "competition law" and "antitrust law" will be used indistinctly throughout this article.

[4] The "hype" surrounding AI has loomed over the field of competition law for some time now. See, for instance, Nicolas Petit, *Antitrust and Artificial Intelligence: A Research Agenda*, 8 J. COMPETITION L. & PRAC. 6, 361 (2017) (noting that the "hype" surrounding AI had reached the antitrust community in 2016).

[5] The many contributions in this book are a prime example.

[6] *See*, *e.g.*, CPI, *AI Boom to Fuel Anticompetitive Behavior in Big Tech, Warns German Antitrust Chief*, PYMNTS (Jun. 26, 2024) https://www.pymnts.com/cpi-posts/ai-boom-to-fuel-anticompetitive-behavior-in-big-tech-warns-german-antitrust-chief/; Speech by EVP Margrethe Vestager at the European Commission Workshop on "Competition in Virtual Worlds and Generative AI", EUROPEAN COMMISSION SPEECH (Jun. 28, 2024) https://ec.europa.eu/commission/presscorner/detail/en/speech_24_3550 (arguing that "now is the time to act");

Indeed, much of this is born out of antipathy to the growth of the internet ecosystem in general, and a fear that the "harm" that arose from the growth of early internet companies will recur. *See, e.g.,* Rana Foorhar, *The Great US-Europe Antitrust Divide*, FINANCIAL TIMES (Feb. 5, 2024), https://www.ft.com/content/065a2f93-dc1e-410c-ba9d-73c930cedc14 (FTC Chair Lina Khan opining that "we are still reeling from the concentration that resulted from Web 2.0, and we don't want to repeat the missteps of the past with AI.")

[7] The concerns of competition authorities over AI are well summarized in an op-ed by *The Economist*. "Broadly speaking, the authorities have two areas of concern. The first is whether the world's biggest companies are trying to tie businesses into their products in anticompetitive ways. The second is about control: are some of the largest generative-AI investments poorly disguised acquisitions intended to sidestep antitrust consideration?" Schumpeter, *Is Artificial Intelligence Making Big Tech too Big?*, THE ECONOMIST (Jun. 23, 2024) https://www.economist.com/business/2024/06/23/is-artificial-intelligence-making-big-tech-too-big.

One difficulty with this narrative, however, is that, strictly speaking, there is no such thing as an "AI market"[8] because AI is not a unitary, monolithic technology.[9] On the contrary, the "AI Stack" is made up of several layers of technology with vastly different features that result in an almost infinite range of potential business and end-user functionalities. Granted, those calling for stringent preemptive antitrust intervention in AI are not oblivious to this fact. For example, Andreas Mundt, the head of Germany's Bundeskartellamt, worried that there is a "great danger that we'll see an even deeper concentration of digital markets and power increases *at various levels*, from chips to the front end where users interact with tech platforms" (emphasis added).[10] Similarly, Margrethe Vestager, the EU's Competition Commissioner differentiated between foundational models and the rest of the value chain.[11] Indeed, in other regulatory contexts, lawmakers often differentiate between so-called "high risk" and "limited risk" AIs,[12] such as between AI techniques used to guide weapons systems or predictive policing on the one hand, and those used to power consumer products like search engines or cars, on the other. While this categorization is crude, it demonstrates that lawmakers are aware that "AI" as a regulatory concept needs to be decomposed into categories that are useful in the real world.

However, and despite these caveats, the full complexity of AI still eludes the antitrust community. Some — maybe even most — of these limitations stem from the habitual challenge of keeping theoretical and conceptual frameworks up to date with brisk technological developments. This is especially apposite in the field of AI, however, which has blossomed at break-neck speed in a variety of directions and across a plethora of fields. For example, as Jonathan Barnett points out in another contribution to this book, there is still no settled definition of "foundation model."[13] The European Commission's 2021 proposal for an AI Act did not even cover foundation models,[14] which shows just how quickly the market is evolving. Our understanding of AI is bounded in other fundamental ways, such as by the lack of clarity surrounding the relation between human cognition and the cognitive tasks that AI is increasingly capable of performing.[15]

---

[8] Impact Assessment Accompanying The Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts {COM(2021) 206 final}, 27 ("This risk is significantly higher when the *AI market* is fragmented with individual Member States taking unilateral actions" (emphasis added)).

[9] *See, e.g.*, Benedict Evans, *The Problems of AI Ethics*, BENEDICT EVANS (Mar. 23, 2024) https://www.ben-evans.com/benedictevans/2024/3/23/the-problem-of-ai-ethics-and-laws-about-ai (arguing that it is pointless to think about the ethics of AI or "regulating AI" because AI is not one thing, but a changing bundle of technology with many different uses which raise a plethora of different questions).

[10] CPI, *supra* n. 6

[11] Vestager, *supra* n. 6.

[12] *See, e.g.*, Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts {SEC(2021) 167 final} - {SWD(2021) 84 final}, at Article 6.

[13] Jonathan M. Barnett, *The Case Against Preemptive Antitrust in the Generative Artificial Intelligence Ecosystem*, in A. Abbott and T. Schrepel (eds.), Artificial Intelligence and Competition Policy, Concurrences, n. 9 (2024).

[14] Thibault Schrepel, *Decoding the AI Act: A Critical Guide for Competition Experts*, Amsterdam Law & Technology Working Paper Series 2024, 17 (2024), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4609947.

[15] *See, e.g.*, Oren Etzioni, *AI's progress isn't the same as creating human intelligence in machines*, MIT TECHNOLOGY REVIEW (June 28, 2022), https://www.technologyreview.com/2022/06/28/1054270/2022-innovators-ai-robots/. That is despite claims that AGI is on the horizon, *see* Leopold Aschenbrenner, *Situational Awareness: The Decade Ahead* (2024),

Much of this ignorance stems from the fact that we do not know that much about how the human mind works, in the first place — let alone AI.[16]

It is naïve, some would even say hubristic, to expect antitrust law to solve most — or even *any* — of the major theoretical conundrums surrounding AI. As a mostly reactive discipline that functions through piecemeal doctrinal progression and the incorporation of mainstream insights from other fields — so far mostly from economics and industrial organization — the best antitrust law can hope to do is observe, learn, and adapt. The trillion-dollar question is how.

In this chapter, we argue that the first step to ensuring that antitrust law stays relevant in the age of AI is developing a principled approach to defining AI relevant markets; one that is legally, economically, and technologically sound. Relevant market definition is central to antitrust law because it is the starting point of most, if not all antitrust law cases.[17] A relevant product market is typically comprised of all those products which consumers view as substitutable and which can therefore be said to compete against each other. By delineating the boundaries of competition between firms, relevant product market definition alerts of the presence of market power and, by extension, of the likelihood of anticompetitive effects.[18] Despite its limitations and despite not being an end in itself,[19] relevant product market definition is, and likely will remain, the

---

https://situational-awareness.ai, the mechanisms of what makes humans "generally intelligent" still elude us. Thus, the very nature of the task that current AIs perform similarly eludes a functional definition of how it actually related to human-level intelligence. It is entirely possible that the cognitive tasks we are modeling are as to full human consciousness as having a thumb is to being an artist. The artist may employ his hands to make a beautiful work of art, or may direct automated tools, but the art is not in the thumb or the tool but in the direction of those tools. Similarly, we do not yet know whether the "thinking" the machines do is more like a tool or is somehow the first step toward consciousness, but it is entirely possible that these systems remain highly sophisticated tools with no internal consciousness.

[16] *Researchers Are Figuring Out How Large Language Models Work*, THE ECONOMIST (Jul. 11, 2024) https://www.economist.com/science-and-technology/2024/07/11/researchers-are-figuring-out-how-large-language-models-work ("Because LLMs are not explicitly programmed, nobody is entirely sure why they have such extraordinary abilities…LLMs really are black boxes."); see also ANIL ANANTHASWAMY, WHY MACHINES LEARN: THE ELEGANT MATH BEHIND MODERN AI 7-25 (2024) (Describing how, despite advances in the ability of machines to recognize patterns, the basic learning mechanisms of even simple biological creatures continue to elude AI researchers).

[17] In this chapter, we focus exclusively on relevant product markets. We do not discuss relevant geographic markets or relevant temporal markets.

[18] Jonathan B. Baker, *Market Definition: An Analytical Overview*, 74 ANTITRUST L.J. 1, 129 (2007)("Throughout the history of U.S. antitrust litigation, the outcome of most cases has surely turned on market definition than on any other substantive issue. Market definition is often the most critical step in evaluating market power and determining whether business conduct has or likely will have anticompetitive effects."); *see, e.g.,* in the US*, Eastman Kodak Co. v. Image Technical Servs. Inc.,* 504 U.S. 451, 469 n. 15 (1992)("Because market power is often inferred from market share, market definition generally determines the outcome of the case").

[19] *See, e.g.*, Magali Eben, *The Antitrust Market Does Not Exist: Pursuit of Objectivity in a Purposive Process*, 17 J. COMPETITION L. & ECON. 3 586, 567 (2021). It is important to recognize that a dogmatic adherence to relevant market definition could, for instance, discount the pressure exercised by products defined as falling outside. *See, e.g.*, Case 1275-1276/1/12/17 *Pfizer and Flynn Pharma v* CMA [2018] CAT 11, para 119. ("It is fallacious to regard as relevant to the competition analysis only those products defined as falling within the relevant market and to disregard entirely any competitive pressure from those products defined as falling outside it."); *see also* RICHARD WHISH & DAVID BAILEY, COMPETITION LAW (10TH ED.) 24 (2021). ("People must not be seduced by numbers [or relevant market shares] when determining whether a firm has market power").

main tool for thinking about the contours of competition between firms for the foreseeable future.[20]

At the time of writing, however, two interrelated issues stand in the way of a constructive understanding of relevant product market definition in AI. The first is the lack of sophistication in demarcating internal product market boundaries. Currently, the antitrust community seems to treat AI more like the enigmatic black monolith from Stanley Kubrick's *2001: A Space Odyssey* (or Arthur C. Clarke's eponymous novel)[21] than the heterogenous, loosely connected "bundle" of technologies that it is. For example, in a recent joint statement, leading antitrust enforcement authorities identify shared competitive concerns across generative AI foundation models and "AI products."[22] But this is like making blanket statements about potential competitive problems in "food markets" or "technology." It is unclear what this level of generality adds to the conversation, or what can be gleaned from it except that "competition is good" — which can also be said about virtually any market that is not a natural monopoly. Similarly, in her speech, Margrethe Vestager mentioned that, if they were allowed to gain control of vital parts of the AI value chain, large tech companies could foreclose "AI competitors."[23] But competitors to whom, or to what? Does the provider of LLMs compete with a company developing computer vision solutions? In theory, both are active in AI. Take a set of products that are more similar: do autonomous drones compete with self-driving cars? Arguably, the core AI developed in those systems will be similar in many respects, yet the products and relevant consumers groups are highly distinct. Or what about AI developed to help radiologists sort through X-Rays and AI systems that help with protein folding problems in medical research? Both are using very similar core technologies, but for very different purposes and for different types of users.

The discussion is not purely academic, either. Other authors have pointed to how lumping different AI systems under the same regulatory framework can lead to questionable outcomes, such as unwittingly favoring one technology over another and thereby stifling innovation.[24] In a similar vein, a flawed understanding of the limits of the several relevant product markets that make up AI is likely to lead to erroneous inferences about market power and to misconstrue competitive dynamics in the AI value chain — leading to sub-optimal antitrust enforcement.

The second hurdle to an effective antitrust approach to relevant market definition in AI is the misconception that AI does not compete with non-AI technology. This, too, is only partially true,

[20] For an opposite view, *see, e.g.*, Daniel A. Crane, *Market Power Without Market Definition*, 90 NOTRE DAME L.REV. 1, 31 (2014-2015). Ten years after the publication of this paper, however, market definition remains the principal tool for inferring market power. For a defense of relevant market definition in antitrust law given managed expectations concerning its utility and function. *See* Eben, *supra* n. 19.

[21] Arthur C. Clarke, *2001: A Space Odyssey*, Orbit & Abacus (1990).

[22] Joint Statement on Competition in Generative AI Foundation Models and AI Products, European Commission, Federal Trade Commission, Department of Justice, and Competition & Markets Authority (Jul. 24, 2024) https://competition-policy.ec.europa.eu/about/news/joint-statement-competition-generative-ai-foundation-models-and-ai-products-2024-07-23_en.

[23] Vestager, *supra* n. 6.

[24] Schrepel, *supra* n. 14, at 10-11(arguing that the AI Act indirectly sanctions AI systems that are easier to control, thus effect tilting the market in favor of AI systems that behave predictably i.e., deterministic AI systems with low or no randomness, over more "creative" or unpredictable AI systems i.e., nondeterministic AI systems with high randomness. And concluding that "neutrality requires imposing different regulatory burdens on different designs).

and is most likely based on the more general perception that AI is idiosyncratic. However, AI was inspired by the human brain and designed with the specific goal of performing tasks that hitherto required the exertion of human intelligence.[25] While the intelligence of AI is not yet at the level of an average human adult,[26] it can already carry out a range of tasks that have traditionally been performed by humans, such as planning, personal assistance, classification, writing texts, composing songs, producing images, making videos, coding, research, strategizing, negotiating, designing ads — to name just a few. Often, the results are nigh indistinguishable from output generated by human intelligence ("HI").[27] Part of this is arguably because AI acquires "knowledge" in a manner not all too different from the way humans do; that is, through induction, deduction and, above all, the recognition of patterns. In fact, AI has even been known to reproduce human biases, such as uncritically citing papers in the field that are already highly cited or sputtering plausible-sounding gibberish. Just like humans sometimes do.[28]

And yet we often hear from enforcers and commentators that AI is unique. Some sustain that it is uniquely dangerous because it perpetuates falsehoods, fallacies, and stereotypes.[29] But so do humans. Others argue that AI will be able to perform tasks that humans cannot, either because these tasks are qualitatively distinct or because they operate at a scale impossible to replicate by HI. This vision of the distinctiveness of AI might ultimately be more persuasive. One thing is clear, however: treating AI as entirely distinct from tasks performed by humans severely downplays the extent to which the two are substitutable. Analogously to how the view that AI is monolithic obfuscates its heterogeneity, an excessive focus on what makes AI unique — as opposed to what makes it similar to HI and HI-powered tasks — is bound to result in erroneous conclusions about market power, competitive dynamics, monopoly, and, ultimately, thwart socially-optimal antitrust enforcement.

In this chapter, we argue that, in order to overcome these blind-spots, enforcers need to substitute hype and pre-existing biases — whether pro or anti enforcement — for careful, evidence-based

---

[25] Economist, *supra* n. 16. LLMs, for instance, are built using "a technique called deep learning, in which a network of billions of neurons, simulated in software and modelled on the structure of the human brain"; Weijie Zhao, *Inspired but not Mimicking: A Conversation Between Artificial Intelligence and Human Intelligence* NAT. SCI. REV. 9, 3 (2022) ("[AI] is a tool inspired by the human brain and empowered by mathematical and computational methods that can realize multiple intelligent behaviors").

[26] Yann LeCun of Meta and Francois Chollet of Google, two respected AI researchers, have said that current AI systems hardly merit being called "intelligence." *A New Lab and a New Paper Reignite an Old AI Debate*, THE ECONOMIST (Jun. 24, 2024) https://www.economist.com/business/2024/06/27/a-new-lab-and-a-new-paper-reignite-an-old-ai-debate. However, according to a recent paper published by Leopold Aschenbrenner, a former OpenAI employee, AI will be as capable as humans at all intellectual tasks by 2027. Leopold Aschenbrenner. *Situational Awareness: The Decade Ahead*, 7 (2024) https://situational-awareness.ai/wp-content/uploads/2024/06/situationalawareness.pdf.

[27] *See, e.g.*, Sankar Sadasivan et al, *Can AI-generated Text be Reliably Detected?* ARXIV (2024), https://arxiv.org/abs/2303.11156 (finding that detectors of LLM-generated text are largely ineffective).

[28] An early adopter of this was the Corporate Gibberish Generator, *available at* http://www.andrewdavidson.com/gibberish/. What made the software fun and widely disseminated was that it generated phrases that sounded truthful, or at least plausible, despite not meaning anything. This, of course, was a parody of the way businesses sometimes communicate through their spokespeople.

[29] Mekela Panditharatne & Noah Giansiracusa, *How AI Puts Elections at Risk – And the Needed Safeguards*, BRENNAN CENTER FOR JUSTICE (July 13, 2023), https://www.brennancenter.org/our-work/analysis-opinion/how-ai-puts-elections-risk-and-needed-safeguards.

analysis.[30] Ultimately, we believe that this is a more "future-proof" approach than committing to an enforcement agenda from the outset or regulating AI under blanket rules.[31] We suggest how this can be done by grasping the internal heterogeneity of AI and by understanding what makes AI similar to HI and HI-powered tasks, thus eschewing simplistic narratives about AI's supposed ubiquity and uniqueness that are bound to impede antitrust law from discharging its social role, which is to protect competition for the ultimate benefit of consumers.

The chapter is organized as follows. In Section I, we show that AI is internally heterogenous, which complicates any claims about an "AI market." This narrows relevant AI markets for the purpose of competition law. In Section II, we explain the similarities in how AI and HI acquire and reproduce knowledge. We use this as a basis to contend that some of the functionalities of AI are not as unique as may initially appear and are substitutable for non-AI powered products, services, and inputs (including HI). This expands relevant AI markets for the purpose of competition law. In Section III, we put forward a tentative set of principles to guide relevant market definition in AI, based on the insights from sections I and II. We argue that the task of competition authorities will be to understand the different relevant product markets that comprise AI, on the one hand, and to separate the uses of AI are reasonably substitutable for tasks performed by HI from those that are not, on the other. To achieve this, enforcers should be guided by three questions:

- Who are the consumers?;
- What is the Product?; and
- Does AI Fundamentally transform a comparable product or service?

Section IV concludes.

I. What Makes AI Markets Internally Heterogeneous

AI is not monolithic. To the contrary, the "AI Stack" is made up of several layers of technology with vastly different characteristics that result in an almost infinite range of potential business and end-user functionalities. As the National Security Commission on Artificial Intelligence has observed:

> AI is not a single technology breakthrough… The race for AI supremacy is not like the space race to the moon. AI is not even comparable to a general-purpose technology like electricity. However, what Thomas Edison said of electricity encapsulates the AI future: "It is a field of fields … it holds the secrets which will reorganize the life of the world." Edison's

[30] For example, we would classify blanket claims such as that AI is going to make "all competition problems worse" (CPI, *supra* n. 6) as based on hype rather than evidence. There is no indication of how AI will affect competition, or whether the effect will be negative, neutral, or positive. "AI hype" is not exclusive to competition law, either. See e.g., Anna Cooban, *AI Investment is Booming, How much is Hype?* CNN (Jul. 23, 2023) https://edition.cnn.com/2023/07/23/business/ai-vc-investment-dot-com-bubble/index.html.

[31] The former is what certain competition authorities seem to be doing. See Joint Statement *supra* n. 22 ; with the latter, we are referring to initiatives such as the AI Act. See EU AI Act, *supra* n. 12.

astounding assessment came from humility. All that he discovered was "very little in comparison with the possibilities that appear."[32]

It is overstated to claim that AI is not a general-purpose field, but directionally, the Commission is correct. AI is in fact a diverse collection of different techniques and technologies that are deployed to handle different tasks across many different industries. First, it is important to note that when we talk about AI we are not talking about "general artificial intelligence" (GAI), or what we normally see in science fiction movies. Notwithstanding predictions that GAI is just around the corner,[33] we do not currently have anything operating on that level, but instead have a collection of technologies that depend upon statistical analysis to approximate human-like intelligence.[34]

More practically, understanding how the "AI stack" works at a high level can help us begin to understand the difficulty with defining broad product markets around heterogeneous technologies. Arguably, the first layer to consider is hardware (semiconductors and raw computing hardware) and "XaaS" services. XaaS is an umbrella term for providers of virtualization or abstraction services for on demand access to storage, processing, and a variety of types of software.[35]

Next, there is the data layer. The foundation of any AI system lies in the data it is trained on. Data can be categorized into structured data, which is highly organized and easily searchable in databases (e.g., spreadsheets with rows and columns), and unstructured data, which lacks a predefined format (e.g., text, images, videos).[36] The quality and quantity of this data crucially defines the efficacy for the performance of the ultimate AI systems. Data collection involves gathering relevant information from various sources, while data preparation includes cleaning (removing noise and inconsistencies), transforming (converting data into a usable format), and labeling (tagging data with appropriate labels for supervised learning). As the saying goes, garbage in, garbage out: This initial stage is fundamental because the accuracy and efficacy of an AI system are directly correlated with the quality of its training data. Thus, the techniques not just for gathering data, but curating (or even generating it in the case of synthetic data), is an enterprise unto itself.

Model training, on the other hand, is the process by which AI systems learn from data.[37] Several techniques are employed in this phase, each serving different purposes and applications.

---

[32] Eric Schmidt, et. al, Final Report, National Security Commission on Artificial Intelligence 7 (2021), https://www.dwt.com/-/media/files/blogs/artificial-intelligence-law-advisor/2021/03/nscai-final-report--2021.pdf

[33] *See e.g.* Aschenbrenner, *supra* note 15.

[34] *See generally* Artificial Intelligence And Life in 2030, One Hundred Year Study On Artificial Intelligence (2016), https://arxiv.org/pdf/2211.06318 ; *see also generally* ANANTHASWAMY, *supra*, note 16.

[35] Romit Dey & George Korizis, *How anything-as-a-service (XaaS) can help reinvent business models and transform outcomes across industries*, pwc, https://www.pwc.com/us/en/services/consulting/business-transformation/library/use-xaas-to-reinvent-business-models.html (*last visited* July 31, 2024).

[36] *See, e.g., Structured vs unstructured data*, IBM, https://www.ibm.com/think/topics/structured-vs-unstructured-data (last visited Aug. 11, 2024); Dongdong Zhang, et al., *Combining structured and unstructured data for predictive models: a deep learning approach*, BMC Med Inform Decis Mak 20, 280 (2020), https://link.springer.com/article/10.1186/s12911-020-01297-6 (describing generally the use of both structured and unstructured data in predictive models for health care).

[37] ANANTHASWAMY, *supra,* note 16 at 12.

Supervised learning involves training the model on labeled data, where the input-output pairs are known, and are used for tasks like classification and regression.[38] By contrast, unsupervised learning involves training on unlabeled data, where the model must identify patterns and relationships within the data itself, commonly used for clustering and dimensionality reduction tasks.[39] Reinforcement learning stands out as a method where models learn through trial and error, making sequential decisions and receiving rewards or penalties based on their actions.[40] This is particularly effective for dynamic and complex decision-making environments.[41] Transfer learning, another critical approach, involves adapting a pre-trained model on one task to perform a related task, significantly reducing the resources required compared to training a model from scratch.[42] Once again, training of models is an enterprise unto itself, involving a variety of hardware components[43] and software components[44] and firms with different levels of specialization in each of these tasks.

Once trained, AI models must be deployed in an environment where they can operate and provide value. The deployment phase can occur in various environments, each with its own set of advantages and considerations. Cloud deployment involves hosting models on cloud platforms, providing scalability and easy access to computational resources, making it ideal for handling large-scale data and serving a global user base.[45] Edge deployment, on the other hand, involves placing models on local devices or edge servers closer to the data source, reducing latency and bandwidth usage, which is crucial for real-time applications like autonomous vehicles and IoT devices.[46] Finally, on-premises deployment entails hosting models on an organization's internal servers, offering greater control over data security and compliance, particularly important in industries with stringent data protection regulations such as healthcare and finance.[47] Each of these deployment methods is managed by different firms operating under different business models. For instance, ChatGPT deploys its service on cloud infrastructure but only allows consumers to access their trained models through a subscription front end. Hugging

---

[38] *See Id.* at 12-13.

[39] *See Id.* at 18-25.

[40] *Id.*

[41] *See Id.* at 24-25 (Describing how unsupervised learning models excel at finding solutions to multi-variate problems involving linear relations, such as learning how to predict housing prices based on a training set with many data points).

[42] *What is Transfer Learning?*, AWS, https://aws.amazon.com/what-is/transfer-learning/ (*last visited* July 31, 2024).

[43] *See e.g., TPUs vs. GPUs: What's the Difference?*, PURESTORAGE (May 2, 2024) https://blog.purestorage.com/purely-educational/tpus-vs-gpus-whats-the-difference/ (last visited July 31, 2024); Josh Schneider & Ian Smalley, *What is a field programmable gate array (FPGA)?*, IBM (May 8, 2024) https://www.ibm.com/think/topics/field-programmable-gate-arrays.

[44] *See e.g. Why TensorFlow*, TENSORFLOW, https://www.tensorflow.org/about (last visited July 31, 2024); *Learn the Basics,* PYTORCH, https://pytorch.org/tutorials/beginner/basics/intro.html (last visited July 31, 2024); *About Keras 3*, Keras, https://keras.io/about/ (last visited July 31, 2024). These software packages are often developed as a complex mix of open-source and proprietary efforts that cross-inform each other.

[45] *See, e.g. Amazon SageMaker*, AWS, https://aws.amazon.com/sagemaker/ (last visited July 31, 2024).

[46] *What is edge AI?*, IBM, https://www.ibm.com/topics/edge-ai (last visited July 24, 2024).

[47] *See e.g. NVIDIA ChatRTX*, NVIDIA, https://www.nvidia.com/en-us/ai-on-rtx/chatrtx/ (last visited July 31, 2024).

Face is an open source project that hosts a wide variety of models, and allows individuals to deploy them in any environment, including either locally[48]or online in a cloud [4950]

Beyond the stack, there is heterogeneity in how AI is deployed and employed. While Large LLMs have garnered significant attention for their capabilities in natural language processing, it is crucial to recognize that the field of AI encompasses a multitude of other technologies, each excelling in different domains. Computer vision models, for example, utilize Convolutional Neural Networks to interpret and analyze visual data.[51] These models are indispensable in applications like medical imaging, where they assist radiologists in detecting abnormalities in X-rays and MRIs, and in autonomous vehicles, where they enable the car's system to recognize and respond to road conditions and obstacles.[52]

As noted above, reinforcement learning is another pivotal AI technology, distinct in its approach to training models through a system of rewards and penalties. This technique is particularly effective in environments requiring sequential decision-making and adaptability. For instance, reinforcement learning has been instrumental in developing advanced robotics, where machines learn to perform complex tasks by optimizing their actions based on continuous feedback. A prominent example is AlphaGo, which mastered the game of Go, a strategy board game, by learning from millions of games and refining its strategy through self-play.[53] Similarly, reinforcement learning is being applied in dynamic resource management and real-time strategy games, showcasing its versatility and effectiveness in scenarios where adaptability and strategic planning are paramount.[54]

Moreover, Graph Neural Networks (GNNs) represent a burgeoning area of AI that excels in handling data structured as graphs.[55] These models are particularly adept at capturing relationships and interactions within complex networks, making them invaluable in fields such as social network analysis, molecular biology, and recommendation systems.[56] For example, GNNs can analyze social media connections to identify influential users or detect communities,

---

[48] *See Stable Diffusion Web UI*, https://github.com/AUTOMATIC1111/stable-diffusion-webui (last visited July 31, 2024).

[50] You can, for example, run a number of models from Hugging Face in Google Collab notebooks. *See Transformers Notebooks*, Hugging Face, https://huggingface.co/docs/transformers/en/notebooks (last visited July 31, 2024)

[51] *Convolutional Neural Networks (CNNs), Deep Learning, and Computer Vision*, intel, https://www.intel.com/content/www/us/en/internet-of-things/computer-vision/convolutional-neural-networks.html (last visited July 31, 2024)

[52] *See* Georgios Kourounis, et. al, Com*puter image analysis with artificial intelligence: a practical introduction to convolutional neural networks for medical professionals*, 99 POSTGRADUATE MEDICAL J. 1178 (2023).

[53] *See Alpha Go*, Google DeepMind, https://deepmind.google/technologies/alphago/ (last visited July 31, 2024)

[54] *See* Ying Chen, et. al, *Deep Reinforcement Learning-Based Dynamic Resource Management for Mobile Edge Computing in Industrial Internet of Things*, 17 IEEE TRANSACTIONS ON INDUSTRIAL INFORMATICS 4925 (2021), *available at* https://ieeexplore.ieee.org/document/9214878 ; *see also* Harshit Sethy, Amit Patel, and Vineet Padmanabhan, *Real Time Strategy Games: A Reinforcement Learning Approach*, 54 PROCEDIA COMPUTER SCIENCE 257 (2015), *available at* https://www.sciencedirect.com/science/article/pii/S187705091501354X

[55] Rick Merritt, *What Are Graph Neural Networks?*, NVIDIA (Oct. 24, 2022) https://blogs.nvidia.com/blog/what-are-graph-neural-networks/

[56] *Id.*

and in molecular biology, they can predict the properties of molecules based on their structural relationships, aiding in drug discovery.[57]

II. What Makes AI Similar to Non-AI Products and Services

In early history, the humanoid giant Talos of Greek mythology and the Golem of Jewish folklore exemplified one of humanity's longest standing fantasies: a helper capable of performing tasks with the same proficiency as a human.

Today, AI can carry out tasks which have typically required the exertion of HI,[58] to the extent that it is often very difficult to tell if output that has been generated by AI or by "the wonderful computers in our head."[59] According to one paper, detectors of LLM-generated text are woefully incapable of detecting AI-generated content.[60] In turn, universities and other institutions are putting in place onerous disclosure rules about the use of AI; tacitly admitting that detection is impossible — or prohibitively complicated — unless disclosed voluntarily.[61] One possible explanation is that AI and HI-generated output are so similar because the two systems acquire and utilize knowledge in comparable ways. As John Searle famously said, "if you can exactly duplicate the causes, you could duplicate the effects."[62] However, the question of what humans know or can know is a difficult one, a dispute that is further complicated by the introduction of AI. Epistemology has been one of the four pillars of philosophy since the time of the ancient Greeks,[63] and it is unlikely that we will contribute to the debate in any significant way here. Suffice it to say that there are epistemological theories which were (obviously) developed with HI in mind but which nevertheless also appear to fit AI, at least to some extent. According to Kissinger, Schmidt and Huttenlicher, for example, AI is more Wittgenstein than Plato.[64] By this they mean that the "knowledge" of AI results from observing the particular qualities of things.[65]

---

[57] *Id.*

[58] One definition of AI is "the ability of software to perform tasks that traditionally require human intelligence." See M. Chui, E. Hazan, R. Roberts et al, *The Economic Potential of Generative AI: The Next Productivity Frontier* MCKINSEY COMPANY 3 (2023) http://dln.jaipuria.ac.in:8080/jspui/bitstream/123456789/14313/1/The-economic-potential-of-generative-ai-the-next-productivity-frontier.pdf.

[59] At Least 10% of Research May Already be Co-Authored by AI, ECONOMIST (Jun. 26, 2024) https://www.economist.com/science-and-technology/2024/06/26/at-least-10-of-research-may-already-be-co-authored-by-ai ("academic policies on llm use are in flux. Some journals ban it outright. Others have changed their minds. Up until November 2023, Science labelled all llm text as plagiarism, saying: "Ultimately the product must come from—and be expressed by—the wonderful computers in our heads").

[60] Sadasivan et al, *supra* n. 27.

[61] *See, e.g.*, European Commission, Living *Guidelines on the Responsible Use of Generative AI in Research*, ERA FORUM STAKEHOLDERS DOCUMENT (2024), 6 (recommending that researchers detail the generative ai tools used in the research process, and how these tools have been used); IE University, *Guidelines for Faculty Use of AI Tools for Academic Work* (2024) https://sites.google.com/view/teachingwithai/ie-statement-policies#h.vl5yg05kvnmq (suggesting that students disclose a list of prompts used and how they were used in producing the relevant output).

[62] John Searle *Minds, Brains, and Programs* BEHAV. BRAIN SCI. 3 417, 422 (1980).

[63] The other three are logic, ethics, and metaphysics.

[64] H. KISSINGER, E. SCHMIDT & D. HUTTENLOCHER, THE AGE OF AI 59 (2021).

[65] As indicated earlier, this is a complex debate. But, oversimplifying, according to Platonic epistemology, humans perceive imperfect reflections of perfect things that exist in the realm of ideas. knowledge is formed by grouping

But why does it matter how AI and HI "think" for the purpose of relevant product market definition? Is relevant product market definition not about product, rather than process substitutability — so that the only thing that matters is that two products are interchangeable, regardless of how they are made? Mostly, yes.

From the perspective of the consumer — which generally also includes the intermediary or business customer under competition law[66] — two systems that think or operate similarly could reasonably be expected to perform the same, or similar tasks, and thus be substitutable. When a potential buyer is deciding whether to hire a human creative or invest in an AI powered generator of text, for instance, he or she will want to know what each one can do. Since it is impossible to test for the full universe of output which that buyer (or employer) will require in the future, nor is it likely that the buyer even knows what those needs will be, a useful heuristic is to know what the AI or HI is capable of doing. Part of this is gauging how they think. Thus, an employer evaluates a candidate in an interview based on their credentials (the knowledge that has been "put" into him or her), portfolio (previous examples of output), and an interview — which is essentially an intelligence or competence test where the employer tries to estimate how successful a candidate is likely to be at resolving current or future work-related problems. The same applies to AI. The potential buyer might want to know what input the LLM has been "trained" on, would want to see examples of previous output (perhaps some texts, lyrics, or articles — depending on the "job"), and would want to test the AI's *ad hoc* response to a series of prompts.

Does this mean that an understanding of AI and HI epistemology is enough to inform relevant product market definition in AI? No. It is only one tool for understanding substitutability.

---

observed phenomena under pre-existing "ideal" or "perfect" categories. For instance, no one has ever seen a perfect circle or a perfect line, yet everyone knows what a perfect circle and a perfect line are. See PLATO, CRATYLUS, para. 389. The early enlightenment refined these ideas further, with a continued focus on classifying observable phenomena according to mechanistic, rational rules. Kissinger, Schmidt & Huttenlocher, *supra* n. 64, at 61. Kant argued in *The Critique of Pure Reason* that the mind structures incoming sensory "data" according to certain impositions that exist *a priori* (before experience) and are thus metaphysical. Earlier, Descartes had rejected the notion that *a posteriori* claims based on observation could be the basis for knowledge because we did not know whether our perceptions were accurate. From the maxim of *cogito ergo sum* — the only principle that could be ascertained as absolutely true — he concluded that material things were knowable not based on sensorial experience, but on their substance. Contrast this with Hume, who broke away from Plato's "substance." *See* DAVID HUME, A TREATISE ON HUMAN NATURE, Book I, Part I, Sect. VI (1739). In contrast to Plato's rationalism (i.e., the idea that knowledge results from absolute, immutable principles that are not learned through experience but implicit in reasoning), Hume argued that knowledge resulted from the observance of similarities. For instance, gold, Hume argued, was just the collection of certain ideas of color, weight, malleableness, fusibility, etc. *Id*. Even the concept of self, Hume contended, is a "heap or collection of different perceptions united together by certain relations and suppos'd, tho' falsely, to be endow'd with a perfect simplicity or identity." *Id*. Part IV, Sect. II. Hume's epistemology laid the basis for successive philosophers to argue, like Wittgenstein did, that knowledge resulted from generalizations about similarities across phenomena, which he termed "family resemblances." Kissinger, Schmidt & Huttenlocher, *supra* n. 64, at 49; see also, generally, LUDWIG WITTGENSTEIN, PHILOSOPHICAL INVESTIGATIONS (2009); *See* also Kissinger, Schmidt & Huttenlocher, *supra* n. 64, at 48-49: "In the late twentieth century and early twenty-first, this thinking informed theories of AI and machine learning. Such theories posited that AI's potential lay partly in its ability to scan large data sets to learn types and patterns — e.g., groupings of words often found together, or features most often present in an image when the image was of a cat — and then to make sense of reality by identifying networks of similarities…even if AI would never know something in the way a human mind could, an accumulation of matches with the patterns of reality could approximate and sometimes exceed the performance of human perception and reason."

[66] Case C-377/20 *Servizio Elettrico Nazionale and Others* ECLI:EU:C:2022:379, para 46; In the EU and UK context, see John Vickers, Competition Policy and the Consumer Welfare Standard, JOURNAL OF ANTITRUST ENFORCEMENT 1, 3-4 (2024) ("In line with the jurisprudence I will take it that [business customers are 'consumers']")

Furthermore, not all AIs are trained in the same way,[67] which can have implications on what they can do and how they can do it. The point here is to underscore that there are overlaps between AI and HI which can also help explain why AI and HI can be used to perform the same, or similar tasks, and why the final output is sometimes so difficult to distinguish. Epistemological considerations can also serve to predict and understand the extent of the substitutability between the two. Ultimately, however, the question of whether there is substitutability will depend on the objective characteristics of the product and, above all, whether consumers *view* them as interchangeable.

And, sure enough, AI can today perform many tasks which would appear to be substitutable for their HI-powered counterparts in terms of style, performance, and quality.[68] As Feurriegel et al put it:

> For a long time in history, it has been the prevailing assumption that artistic, creative tasks such as writing poems, creating software, designing fashion, and composing songs could only be performed by humans. This assumption has changed drastically with recent advances in [AI] that can generate new content in ways that cannot be distinguished anymore from human craftsmanship.[69]

"Creativity" is no longer a human-only endeavor.[70] There are hundreds of GAI applications that are able to produce text (ChatGPT), images (Dall-E), speech (Speechify), music (Suno), code (GitHub Copilot), and video (Runway). Some of these are generalist ("horizontal), while others are specialized ("vertical"). For example, ChatGPT is a generalist multimodal model ideal for text and image generation, while Reword specializes in writing blogs and SocialBee specializes in social media posts. Other specific applications include marketing, innovation management, scholarly research, and education.[71] Generative AI can "increase efficiency and productivity by automating many tasks that were previously performed by humans, such as content creation, customer service, code generation, etc."[72] What a GAI application is "good" at will depend on the FM and input it has been trained with. Thus, some text generators will be better at producing one type of output than another – even within a specialized category – depending on their training and subsequent "fine-tuning."[73] The possibilities borne of different combinations of training, fine-tuning, and the underlying technology used are almost endless. In addition to creating new content based on learned patterns, AI can also be used to assist humans such as, for instance, by checking style, punctuation, grammar, clarity, engagement, and spelling

[67] Stefan Feuerriegel et al, *Generative AI*, Bus Inf Syst Eng 66, 111 (2024).

[68] *Id.* at 166 (stating that generative AI is a "human-task technology.").

[69] Feuerriegel et al, *supra* n. 67~~67~~, at 111.

[70] For a discussion on whether there even such a thing as creativity is, see Minsky, 1982: 5. (Arguing that creativity is just a shorthand for better-knitted, ordinary virtues. The upshot is that machines can at least be creative in the same way as humans. In other words, if machines cannot be "creative," neither can humans); see also Feurriegel et al. at 116. (Arguing that while AI was in the past mostly understood to be analytic, whereas now AI has gained the capability to perform generative tasks suitable for content creation by combining elements in novel ways.)

[71] Feurriegel et al. *supra* n. 67~~67~~, at 111, and the literature cited therein.

[72] *Id.* at 120.

[73] *Id.* at 117. (Arguing that the correctness of generative AI models is highly dependent on the quality of training data and the according leaning process.)

Formatted: Spanish (Spain)

Field Code Changed

Formatted: Spanish (Spain)

Field Code Changed

(Grammarly). If it were human, Grammarly would be an editor. A recent paper showed that a personalized GPT-4 powered model was 82% more persuasive than a human in debates, suggesting that AI could also play the part of the debater (or sophist?).[74] AI can also serve to improve not just human output, but input (skills) by exposing HI to new ways of thinking. For instance, [75]

All of the tasks outlined above can and still are performed by humans; by a research assistant, a clerical assistant,[76] a musician, a translator, a singer-songwriter, software developer, editor, coach, a blogger, etc. The list of tasks a, say, foundation model can perform is vast.[77] In time, generative AI applications will likely transition from dispatching mundane tasks to more sophisticated ones, including passing moral judgment[78] (assuming this is not already the case).[79] As a result, the line between HI and AI output is likely to become even more blurred. Some studies even estimate that AI could replace 300 million jobs of knowledge workers, with 900 occupations exposed to some degree of automation by AI (most jobs, however, are likely to be complemented, rather than substituted by AI, the study points out).[80]

But let us not understate the counterfactual, either. We have so far talked about instances in which AI and HI might be similar. But there are also important differences between AI and HI.[81] On the one hand, AI can "produce phenomena that are truly new, not simply more powerful or efficient versions of things past."[82] This could come about, for instance, through a qualitatively different functionality that cannot be replicated by human capabilities at all, such as a way of processing data that is fundamentally alien to humans. Or it could come about as a result of compressing the time frame in which results are achieved to an extent that renders tasks

[74] Francesco Salvi et al., *On the Conversational Persuasiveness of Large Language Models: A Randomized Controlled Trial*. Working Paper (2024), https://arxiv.org/pdf/2403.14380.

[75] Shin et al., *Superhuman Artificial Intelligence Can Improve Human Decision-Making by Increasing Novelty,* 120 PROC. NAT'L ACAD. SCI. 12 (2023).

[76] One study in fact finds that it is this group that will be the most exposed to automation from generative AI. Paweł Gmyrek, Janine Berg & David Bescond, *Generative AI and jobs: A global analysis of potential effects on job quantity and quality*, ILO Working Paper 96 (2023).

[77] *See* Thibault Schrepel & Alex Pentland, *Competition Between AI Foundation Models: Dynamics and Policy Recommendations*, Amsterdam Law & Technology Institute (ALTI) Working Paper 3-2023, at n. 2 (2023). Also, at 1, arguing that "recent advances in deep learning have given rise to foundation models that underpin an infinite number of generative AI applications."

[78] Feuerriegel et al, *supra* n. 67~~67~~, at 116.

[79] Eyal Aharoni et al., *Attributions Toward Artificial Intelligence Agents in Modified Turing Test*, SCI. REP. 14 (2024) (Finding that people rated AI's moral reasoning as superior in quality to humans' along almost all dimensions, including virtuousness, intelligence, and trustworthiness).

[80] Joseph Briggs & Devesh Kodnani, *The Potentially Large Effects of Artificial Intelligence on Economic Growth*, GOLDMAN SACHS, 1 (2023). *See also id.* at 7, stating that "our estimates intuitively suggest that [...] 18% of work globally could be automated by AI on an employment-weighted basis."; *see also* Xiang Hui, Oren Reshef & Loufeng Zhou, *The Short-term Effects of Generative Artificial Intelligence on Employment: Evidence from the Online Labor Market*, CESIFO WORKING PAPERS (2023)(finding that "freelancers in highly affected occupations suffer from the introduction of generative AI, experiencing reductions in both employment and earnings," and that "in the short term generative AI reduces overall demand for knowledge workers of all types").

[81] See, in general, Johan Egbert Korteling et al., *Human- versus Artificial Intelligence*, FRONT. ARTIF. INTELL. (2021); Zhao, *supra* n.25, at 4.

[82] Hazan, *supra* n. 58, at 50.

Formatted: Spanish (Spain)

Field Code Changed

performed by humans obsolete, or relegates them to a different category (e.g., AI's fast data processing might be able to automate output personalization across various product categories with a speed that is impossible to replicate by HI, thus creating a new type of product altogether).[83] Indeed, signals from AI systems propagate much quicker than human nerves, which operate at the speed of at most 120/ms.[84] Similarly, the amount of cognitive information we can retain and our ability to "multi-task" is severely limited.[85] Ultimately a sufficiently large quantitative improvement can become indistinguishable from a qualitative leap. Furthermore, an AI-enabled reduction in production costs could lead to new capabilities, not just increased output.

For instance, according to one study, robotic process automatization will not only improve handcrafted processing rules but "enable entirely new types of automatization by retrofitting and thus intelligentizing legacy software."[86] The authors anticipate the development of a new generation of process guidance systems in business process management:

> While traditional system designs are based on static and manually-crafted knowledge bases, more dynamic and adaptive systems are feasible on the basis of large enterprise-wide trained language models (internal citations omitted).[87]

By the same token, it could be that AI is incapable (or incomparably worse) at performing tasks that can be easily carried out by HI, thus limiting substitutability in exactly the opposite direction. In 1988, Moravec wrote:

> It is comparatively easy to make computer models exhibit adult level performance on intelligence tests or playing checkers, and difficult or impossible to give them the skills of a one-year old when it comes to perception and mobility.[88]

As Korteling et al have pointed out, "Moravec's Paradox implies that HI and AI are intelligent in different ways."[89] Generally, people are better at carrying out a broader spectrum of cognitive and social tasks under uncertainty.[90] Recent developments in image recognition however challenge this view. Foundation models especially have been trained to recognize images with deep learning technology, "which is based on some principles of biological neural networks."[91]

---

[83] Think, for example, of advertising, summarizing web pages for mobile devices, creating songs or playlists, generating relevant instant social media feeds, or even adapting political narratives instantaneously depending on the audience. These are all things that HI is capable of, it just generally requires more time. In addition, humans can retain less information that AI due to memory constraints.

[84] Korteling et al., *supra* note 81, at 4-5 (2021). But note, for instance, that human brains are millions of times more efficient in energy consumption than computers. *Id.*

[85] *Id.* at 4.

[86] Feurriegel et al., *supra* n. 67~~67~~, at 118, and the papers cited therein.

[87] *Id.*

[88] Hans Moravec, MIND CHILDREN (1988).

[89] Korteling et al., *supra*, n. 84~~84~~, at 6. *See also* Anathaswamy's description of the intelligence of a duckling, and how even that basic level of automatic learning continues to elude computer scientists. ANANTHASWAMY, *supra* note 16, at 7-8.

[90] *Id.* at 7.

[91] *Id.*

Formatted: Spanish (Spain)

Field Code Changed

As a result, one of the major "weaknesses" (or differences) of AI compared to HI, i.e., the difficulty of interpreting human language, symbolism and context — which typically requires an extensive frame of reference[92] — is potentially being bridged.

There is no hard and fast rule, and substitutability will need to be studied on a case-by-case basis for every AI product or service. But, at the most general level, the question that competition law will have to answer is which type of AI are we dealing with? The type that is reasonably substitutable for human capabilities, or the type that is not? In the end, this is what is going to inform the boundaries of relevant product market definition.

III. Tentative Principles for Market Definition in AI

Jonathan Barnett has observed that in the GAI market, a preemptive approach to antitrust is dangerous to consumer welfare:

> At the early stages of a market's development, uncertainty concerning the competitive effects of certain business practices is likely to be especially high, which supports concerns that preemptive intervention would result in significant false-positive error costs by potentially suppressing practices that are either innocuous or yield procompetitive efficiencies. This suppressive effect arises both by constraining existing practices and limiting the future range of transactional innovation. Hence, without grounds to anticipate a future anticompetitive outcome at a sufficiently high level of confidence, there would seem to be a strong presumption against preemptive intervention in the GAI market at its current nascent stage of development. It remains to consider whether there are sufficiently compelling factual or other grounds to overcome that presumption.[93]

Barnett's observation is part of a larger exhortation to enforcers to refrain from jumping to antitrust remedies for feared anticompetitive harms that have yet to materialize. Barnett is correct in general, and specifically if we apply his observations to market definition for AI. Not only are the relevant AI markets nascent, as Barnett notes, but what will even count as a relevant product is as-yet undefined.[94] As we have repeated in various ways: there is no single thing as "AI," even within areas we currently think of as unified. In Barnett's paper, for example, he notes numerous possible divisions within antitrust-relevant markets just for generative AI.[95]

What is needed at the moment, therefore, is not an authoritative definition of what counts as a relevant AI market for antitrust purposes, but a good set of questions that can be asked to help understand when a relevant market emerges.

In competition law, a relevant market is comprised of the relevant product market and the geographic market. In this paper, we have focused on the former. An authoritative competition law textbook defines a relevant product market thus:

---

[92] *Id.* at 7; Zhao, *supra* n.25, at 4.

[93] Barnett, *supra*, n. 13, at 8.

[94] *Id.* at 5.

[95] *Id.* at 3-6.

> The definition of the market is essentially a matter of interchangeability. Where goods or services can be regarded as interchangeable, they are within the same product market.[96]

Accordingly, the definition of the relevant product market hinges on product substitutability. The question that is asked is "would consumers view these two products as substitutable?" If the answer is "yes", then both products form part of the same relevant product market. If the answer is "no", then they do not. The sort of evidence that may be used in defining relevant product markets includes evidence of substitution in the recent past, quantitative tests (e.g., own-price elasticities and cross-price elasticities), views of customers and competitors, market studies and consumer surveys, barriers (including regulatory barriers) and switching costs.[97]

The problem, of course, is that none of this data may be readily available today.[98] The process of market definition in AI is further complicated by the fact that even AI-based companies do not always seem know what the technology is and what it is for.[99] The default fallback marketing trope that AI "is for everything"[100] will, if taken literally (or almost literally), likely to lead to erroneous antitrust outcomes. A more principled and informed approach is needed.

A. Who are the consumers? What is the product or service?

The first, and most basic questions are: who are the consumers and, relatedly, what is the product? In the rapidly evolving field of AI these questions remain elusive. For example, the market for semiconductors implicates advanced processors and chipsets, but the market for those chipsets can have multiple consumers. High-end chips are useful for not just AI, but also for gaming, advanced (non-AI) mathematical modeling, crypto mining, etc. Moreover, even if we try to restrict the market consideration to just AI-relevant consumers, as Barnett notes, there are endemic make-or-buy considerations throughout the AI value chain.[101]A firm can stand up a completely vertically integrated solution, purchasing chipsets, designing hardware, setting up data centers, etc., or it can outsource a number of its operations to external firms, like cloud providers. The market for different types of hardware components consists of a large number of heterogeneous parties with different demand elasticities (indeed, some of the parties view their potential competitors for purchasing hardware as potential partners for providing hosted-access to hardware).

---

[96] WHISH & BAILEY, *supra* n. 19, at 26 (2021); *see also* Case 6/72 *Europemballege Corpn and Continental Can Co Inc v Commission* EU:C:1973:22, para 32; Case T-321/05 *AstraZeneca AB v Commission* EU:T:2010:266, paras 30-31.

[97] *Id.* at 32.

[98] This is an overarching problem in relevant product market definition, not just in the field of AI. As Richard Whish and David Bailey note, in practice, the measurement of interchangeability can be difficult for several reasons: there may not be data on the issue, or the data that is available might be unreliable, incomplete or deficient in some other way. *Id.* at 26.

[99] Benedict Evans, *The AI Summer*, BENEDICT EVANS (Jul. 9, 2024) https://www.ben-evans.com/benedictevans/2024/7/9/the-ai-summer.

[100] *Id.*

[101] Barnett, *supra* n. 13, at 5-6.

Moving further up the stack, there are a number of firms that provide a wide variety of services. For example, firms specialize in providing curated data,[102] testing/red-teaming services,[103] and other types of services that could otherwise be developed in-house as part of a full-stack solution (e.g. providing low-level open-source libraries for use in larger machine learning systems[104]). Generally speaking, the consumers of these products will be other firms developing AI tools, but spread out across a wide array of industries and applications. But, again as Barnett notes, the presence of the make-or-buy decision is relatively complicated in these markets because many of the tools that are currently in use can be deployed internally within a firm using publicly available scientific papers and tools, or they can be purchased where efficiencies are greater.

At the same time, firms can provide the development of foundation models, fine-tuning of models, or both. The consumers here will be even more heterogeneous. Some will be firms using models as modules within their own consumers-facing products or to improve their own processes. Those same models can be deployed to consumers through web interfaces (e.g. ChatGPT, Claude). But the same basic model can power both the firm-facing and the consumer-facing experience, or can be the same model but tailored to the needs of a particular customer on demand. Which leads to the question: is it the model that is forming the relevant market? Or is it the particularly customized model prepared by a particular party that forms the relevant market? The further you get toward restricting the relevant markets based on user-customization, the more you converge on extremely restricted product markets.

At the same time, as you pull back from the customizations of individual consumers the more generalized the model, and the much larger both the consumer pool as well as the pool of potential substitutes. That is to say, if, for argument's sake, we assume a relevant market is, e.g., a customized version of ChatGPT trained on a firm's internal data, the relevant market in that case might be exactly that instance of ChatGPT with exactly one relevant consumer. Hardly an analytically useful market.

Pulled back slightly you could construe the market as the market for model customization services, and pull in possible alternative providers, like Anthropic and Mistral,[105] as well as open-source alternatives. Such a move would bring in a larger amount of potential providers and would be more analytically interesting but then you begin to introduce further complications. What exactly is the service being provided? Is it the training services or the end-consumer product of having some usefully indexed and accessible compendium of their relevant data packaged together with other retrieval services? Which leads to the next question: what exactly is unique about the AI contribution relative to its unique costs.

---

[102] *See, e.g.*, databricks, https://www.databricks.com/ (last visited July 31, 2024); scale, https://scale.com/ (last visited July 31, 2024).

[103] *See, e.g.*, *AI red teaming, protect your AI systems*, Reply, https://www.reply.com/en/cybersecurity/ai-red-teaming (last visited July 31, 2024)

[104] *See, e.g.*, PyTorch and TensorFlow, *supra* note 44.

[105] Mistral is an interesting example because they provide their models open-weight (meaning that anyone can fine-tune them using their own infrastructure or through a cloud provider), but they also provide a fine-tuning service (like OpenAI does for their models).

Pull-back even further to the domain of so-called "foundation models." From the perspective of the customer, these foundation models may not be fully interchangeable, or interchangeable at all. As Thibault Schrepel and Alex Pentland have pointed out:

> Foundation models are commonly observed by policymakers and social scientists at the species level (i.e., "foundation model" as a class), but these lenses fail to see the inherent diversity within the species.[106]

There is no accepted taxonomy of foundation models. The tentative classifications that have been attempted, however, point to how foundation models can have different characteristics with potentially different uses, ultimately serving customers. For example, depending on their training data, foundation models could be general purpose or domain specific.[107] The former are trained on a large variety of data with the aim of performing tasks in all possible domains (e.g., ChatGPT and Google Bard). The latter are designed for a specific task or topic, and are thus trained on more granular, specialized data. If someone is looking to produce a generalist generative AI application, they may not consider domain specific foundation models to be substitutable for general purpose foundation models. Furthermore, someone looking to make a specialized generative AI application on one topic — e.g., finance or cooking — might not find a foundation model trained on specialized data in another field to be interchangeable. The authors make further subdivisions, such as personal foundation models (typically pre-trained on large data sets and fine-tuned on individual's private data) and ecosystem foundation models (trained on data that is not publicly available such as data from different companies in the same industry).[108] These models all have different returns and limits,[109] and they experience different competitive dynamics.[110]

It seems reasonable to assume that these foundation models underpin different end-uses and are thus not fully interchangeable at the intermediary stage, either. Indeed:

> Access to *unique* data sets is critical. There are two reasons for this. First, access to unique datasets may be necessary to provide the specific answer that users of foundation models are looking for. […] Second, these datasets may play a critical role in the overall training of foundation models.[111]

Moreover, inputs and outputs of foundations models vary. Some foundation models are unimodal while others are multimodal. "Unimodal models take instructions from the same input type as their output (e.g. text). On the other hand, multimodal models can take their input from different sources, and generate output in various forms."[112] For example, GPT-4o, which

---

[106] Schrepel and Pentland, *supra* n. 77, at 3.

[107] *Id.*

[108] *Id.* at 4-6

[109] *Id.* at 12.

[110] *Id.* at 23.

[111] *Id.* at 9.

[112] Feuerriegel et al, *supra* n. 67~~67~~, at 113; *see also* Barnett, *supra* n.13, at 3-4.

Formatted: Spanish (Spain)

Field Code Changed

underpins OpenAI's ChatGPT, accepts both image and text output to generate text. MusicLM is text-to-music; AlphaCode is text-to-code, and so on.[113]

Clearly, there are limits to the substitutability between foundation models that produce different output at the modal level, such as text, image, code, or audio. Customers who want images, may not want text, audio or code. But a customer may also not view different inputs as substitutable based, for instance, on the availability of either input. The upshot is that foundation models could, in principle, also be broken down at the modal level depending on their respective input and output data modalities.[114]

But if we take another step, the systems level of generative AI is also quite diverse. Generative AI systems comprise not only generative AI models but also the "underlying infrastructure, user-facing components, and their modality as well as the corresponding data processing (e.g., prompts)."[115] Generative AI systems make the underlying mathematical model "usable" across real-world cases by enabling user interaction through a practical interface.[116] For example, generative AI that produces text often use conversational agents and search engines, image generating AI may use bots, etc. How easy generative AI is to use may be a crucial contributing factor to its success,[117] and thus a relevant competitive parameter vis-à-vis other systems and products.

Other important questions arise across the GAI technology stack. For example, are large datasets substitutable for small datasets? Some suggest that they are.[118] Are open-source and closed-source foundation models substitutable? The answers to these, and similar questions will depend on who the consumers are and what they want the AI product or service for i.e., on the product's function and utility. Evidently, this is not something that can be resolved in the abstract; instead, it will require analysis on a case-by-case basis.

B. Does AI fundamentally transform the product or service?

The third question is going to be the hardest for enforcers to grapple with. The overarching theme is that, if AI furnishes a product or service that is essentially not reasonably replicable by HI, it constitutes a separate product market. Today, AI is shiny and new (at least in the popular imagination) but, with time, we will inevitably become intimately familiar with the limitations of

---

[113] *Id.* at 113.

[114] See also Barnett, *supra* n. 13, at 11. (Arguing that "there is a strong likelihood that the models layer of the GAI ecosystem will devolve into multiple differentiated [foundation model] segments tailored for particular industries or uses…Foundation model markets may disaggregate based on output or input modalities or specific uses.") (internal parenthesis omitted for clarity).

[115] Feuerriegel et al, *supra* n. 67~~67~~, at 113.

[116] *Id.*

[117] For example, ChatGPT's ease of use, especially for non-experts, was a core contributing factor to its worldwide adoption. Feuerriegel et al. *supra* n.67, at 114-115.

[118] Thibault Schrepel, *Alternatives to Data Sharing*, THE REGULATORY REVIEW (Feb. 21 2022) https://www.theregreview.org/2022/02/21/schrepel-alternatives-data-sharing/; Igor Susmelj, *Optimizing Generative AI: The Role of Data Curation*, LIGHTLY (last visited jun. 15, 2024), https://www.lightly.ai/post/optimizing-generative-ai-the-role-of-data-curation (discussing the importance of data curation to ensure that datasets are devoid of noise, irrelevant instances, and duplications, thus maximizing the efficiency of every training iteration).

Formatted: Spanish (Spain)

Field Code Changed

AI-powered products. Indeed, already we are beginning to come to grasp with the awesome power needs of generative AI.[119] The costs for chips are also high and growing as the need for more data centers increases.[120] Cooling costs are also very large for the data centers powering consumer AI applications.[121]

At the moment, many of these costs may be dwarfed by the enthusiasm of investors, but if the scale of the potential power demand is at all in line with reality, very quickly many of these services will need to be more adequately priced into the cost of accessing AI services. This does not mean that AI will go away; to the contrary, we foresee it being a major component in the future economy. What it does mean is that the tradeoffs inherent in using AI for a particular application will become more apparent and will force consumers of all types to evaluate AI not on its own, but in comparison to other possible products and services that, while inferior on some dimensions, will be effective substitutes.

Returning to the example of the custom GPT that a firm wishes to train on its own internal documents. That could be a service operating in a market for "GPT customization services" if we look at it from the provider's perspective. But from the consuming firm's perspective it could be comparing alternative document archival systems, some of which will be based on "good enough" traditional relational databases, some of which on semi-intelligent unstructured retrieval systems, and some on AI-powered LLMs. Arguably, a MySQL-powered traditional application based on well-trod search algorithms is not as powerful as an LLM-powered system that allows you to interactively "chat" with your document history. But the substitutability is based not just on what the more powerful application is but, given the tradeoffs necessary, which products are good enough alternatives. The cost and time involved in standing up an LLM, including coding data to be the most effective for training, and the direct or indirect power and cooling costs of having constantly available LLMs may not make sense when well-understood (but maybe boring) traditional alternatives exist. It is not clear that such traditional applications *should* exist in the same product market as an LLM, but our goal is to merely point out that is it not clear yet that they do not do so.

On the other hand, there might genuinely be AI functionalities that bring entirely novel products to the market. For example, it is possible to imagine that in litigation, AI processing of documents, and interactive discussion with an LLM about the corpus of documents produced in discovery will simply be so much more effective than HI-driven traditional solutions that it does not make sense to treat them as substitutes. Large litigations can be extremely expensive, and, indeed, AI may drive down the marginal cost of document review in a way that it makes it possible for more parties to engage in litigation. We are not passing judgment on whether this is good or bad, but merely observing that such an augmentation offered by AI could be a relevant

---

[119] Beth Kindig, *AI Power Consumption: Rapidly Becoming Mission-Critical*, FORBES, June 20, 2024, https://www.forbes.com/sites/bethkindig/2024/06/20/ai-power-consumption-rapidly-becoming-mission-critical/

[120] Angus Loten, *Rising Data Center Costs Linked to AI Demands*, WALL ST. J. (Jul. 13, 2023) https://www.wsj.com/articles/rising-data-center-costs-linked-to-ai-demands-fc6adc0e

[121] David Berreby, *As Use of A.I. Soars, So Does the Energy and Water It Requires*, YALE ENVIRONMENT 360, Feb. 6, 2024, https://e360.yale.edu/features/artificial-intelligence-climate-energy-emissions

product market in this case with which traditional HI-driver alternatives simply would not be able to compete.

Furthermore, a sufficiently notable quantitative leap can produce a qualitative different product that is not interchangeable for the "old" version. For instance, although HI can also analyze data patterns, AI can analyze vastly bigger swathes of data much more quickly and possibly uncover correlations that would escape the human mind. AI can also work 24/7, without the need for rest, which applies across all tasks. Similarly, AI can perform repetitive tasks without fatigue or degraded performance. Granted, this is also true for traditional software, but in the context of AI-HI substitutability, it might tilt the decision in one direction or the other depending on the extent to which the ability to work non-stop constitutes an important competitive dimension of the product at hand (e.g., it might be decisive in the context of HI vs. AI powered security systems, but less important when it comes to editing blog posts).

Thus, AI can not only create qualitatively new uses but improve existing ones so much that they effectively become a new relevant product market. Competition law recognizes that products belonging to different quality segments can constitute different relevant product markets, similar to how, for example, "premium smartphones" are in a separate category from other smartphones.[122]

IV. Conclusion

Abstraction usually follows the observation of similar phenomena. As such, it is the basis for all deductive reasoning and a fundamental pillar of human epistemology. However, abstractions can sometimes lead to erroneous inferences by ascribing general properties to phenomena that are, in reality, distinct. In the social sciences and in popular discourse, we refer to erroneous abstractions as "generalizations." Obviously, the opposite is also possible; such as when we fetishize small differences and confound novelty — or minor idiosyncrasies — with uniqueness.[123] The same applies to AI. From the outside, AI is an imposing block — not unlike the black monolith from Stanley Kubrick's *2001: A Space Odyssey*. Upon closer examination, however, it turns out that AI is neither homogenous nor as unique as would initially seem.

The lack of a proper understanding of the outward and inward boundaries of AI markets has practical implications for antitrust policy and regulation because it may lead to inaccurate assessments of market concentration and market power, resulting in both under and over-enforcement of competition law compared to the social optimum. For example, it is likely — or at the very least plausible — that as soon as one accounts for the substitutability of AI and non-AI product, the concentration in some of those markets that hitherto appeared to be monopolistic withers away. What changes is not the observed phenomena, but the level of

---

[122] *Justice Department Sues Apple for Monopolizing Smartphone Markets*, U.S. Dept. Justice (Mar. 21, 2024), https://www.justice.gov/opa/pr/justice-department-sues-apple-monopolizing-smartphone-markets.

[123] One manifestation of this is, arguably, Freud's "narcissism of the small differences." *See, generally*, Sigmund Freud, CIVILIZATION AND ITS DISCONTENTS (1929).

abstraction: we understand that the market is more competitive than initially envisioned because competitive pressures is exerted both from within and without AI-specific products.

Conversely, a failure to account for the internal heterogeneity of AI could lead to an under-estimation of market concentration and market power, by artificially expanding the universe of products that comprise the same relevant market. This currently does not seem to be what is happening, however, as the predominant narrative is, somewhat paradoxically, that AI markets are both extremely broad — encompassing a range of different products and technologies — and extremely concentrated. Ultimately, this may be due more to a combination of the lack of sophistication in demarcating AI markets, technological anxiety, pro-enforcement bias, and prejudice against large technological firms, than the articulation of a principled creed. And yet it must be noted that the upshot of this distorted perception, if it ever translated into enforcement, will be a distorted antitrust policy.

But the problem is not just that the over and under-estimation of the relevant product market mystifies the appraisal of market power. Confounding the boundaries of product markets obfuscates the real competitive dynamics in those markets, rendering the enforcer myopic to incentives, the direction of competitive threats, the potential for procompetitive benefits, and the possibility and likelihood of entry and expansion. Such an unprincipled approach would force authorities to rely on tasseography and pre-existing biases to address potential competition problems in a rapidly changing environment that is inherently hostile to such unwavering assumptions. In the world of AI, that is called "hallucination."