

ARTICLE

Fueling Concentration: AI Agents and Network Effects

Raz Agranat & Michal S. Gal

Abstract. AI agents lie at the frontier of the evolving AI economy, functioning as applications that autonomously execute tasks on users' behalf. While contemporary discourse on market concentration in AI ecosystems predominantly focuses on data access asymmetries, this paper examines the susceptibility of AI agent markets to tipping. We posit that traditional network effects pose a significant and underappreciated threat to competition between AI agents.

*

Specifically, the paper warns that as agents increasingly mediate purchasing and selling tasks, they consolidate previously discrete decisions of buyers and sellers, thereby aggregating their bargaining power. This consolidation creates a self-reinforcing positive feedback loop: larger user bases confer augmented negotiation leverage over trading counterparties, enabling AI agents to deliver better outcomes, and attract even more users. Over time, dominant AI agents may increasingly outperform competitors primarily through network size advantages rather than technological superiority. Drawing on network science, the paper illustrates how these dynamics can lead to a scale-free market structure in the burgeoning AI agent economy, with dominant agents forming powerful hubs. It critiques the limits of contemporary antitrust tools in addressing agents' aggregation of market power, particularly when such aggregation transcends traditional market boundaries or circumvents established market-share thresholds.

The paper concludes by emphasizing the rapidly narrowing window for effective regulatory intervention and suggests targeting the preferential attachment

mechanism that fuels the emergence of AI agent hubs as a means of preserving competition in AI agent markets, in large part on the relative importance of long-run learning effects.

1. Introduction

The discourse surrounding artificial intelligence (AI) market concentration has predominantly focused on data accessibility as the fundamental entry barrier. The prevailing theory maintains that data constitutes an essential input for AI services¹ and, accordingly, warns that AI markets are characterized by “data network effects.”² The concern is that a large user base enables the collection of valuable data used for iterative AI training which increases service quality, subsequently attracting additional users in a positive feedback loop.³ Consequently, only a select group of technology enterprises—predominantly large platforms and established AI providers—possess the capacity to obtain the extensive data volumes required to compete effectively in some AI offerings. New entrants, lacking a substantial user base, face significant impediments in acquiring adequate data resources to develop viable competitive models.⁴ While additional barriers—including computational infrastructure, storage capacity, and specialized expertise—constrain new entrant participation in the evolving AI sector,⁵ data accessibility is frequently cited as the most critical barrier and the paramount determinant of competitive dynamics.⁶ The recent rise of GenAI models—especially Large Language Models (LLMs) which consume remarkable amounts of data—has intensified concerns regarding market concentration driven by asymmetrical access to data.

The significance of data as the main entry barrier warrants critical reexamination, at least for some AI offerings. Numerous types of data are not inherently scarce and remain accessible through channels independent of established AI providers or platforms. The near-costless replicability of data implies that AI entrants can effectively utilize both publicly available data or acquire it from data traders.⁷ Furthermore, contemporary computational techniques—including

¹ Dunham, W., 2006, ‘The Determination of Antitrust Liability in *United States v. Microsoft*: The Empirical

² Compare *United States v. Google LLC*, 2024 WL 3647498, at *111-112 (D.D.C., 2024).

³ Id.

⁴ See, e.g. Michal S. Gal & Daniel Rubinfeld, Algorithms, AI and Mergers, 85 Antitrust L. J. 683, 690-691 (2024).

⁵ See Thibault Schrepel & Alex ‘Sandy’ Pentland, Competition between AI Foundation Models: Dynamics and Policy Recommendations, 1 Indus. & Corp. Change 4 (2024), <https://doi.org/10.1093/icc/dtae042> (emphasizing the importance of high quality data in particular).

⁶ Gal & Rubinfeld, *supra* note 4.

⁷ Darren S. Tucker & Hill B. Wellford, Big Mistakes Regarding Big Data, Antitrust Source 7 (2014); Mark MacCarthy, Big Data is Not a Barrier to Entry, Tech. Pol’y Persp. (January 29, 2018),

transfer learning and synthetic data—offer viable alternatives for training AI models that circumvent extensive dataset requirements.⁸ Additionally, some data has only transient value, implying that its utility erodes over time, reducing the advantage of stockpiling it for future purposes.⁹ Empirical evidence also points to diminishing marginal utility of data in powering AI systems, with performance improvements becoming negligible beyond certain thresholds.¹⁰ Even in the context of LLMs, the advent of DeepSeek suggests that GenAI entrants can train their models with less data than initially assumed.¹¹

Against this backdrop, we highlight the role of another barrier in driving market concentration that, while acknowledged in other contexts, remains largely overlooked in analyses of AI markets: traditional (non-data) network effects. This barrier is especially influential in the emerging and vanguard sub-sector of AI markets—AI agents—which has far-reaching implications for the burgeoning AI economy. Drawing on network science, the paper exemplifies how these dynamics can produce a scale-free market structure in the burgeoning AI agent economy, with dominant agents forming powerful hubs. It critiques the limitations of contemporary antitrust tools in addressing agents' aggregation of market power, particularly when such aggregation transcends traditional market boundaries or circumvents established market-share thresholds. The paper concludes by emphasizing the rapidly narrowing window for effective regulatory intervention and suggests targeting the preferential attachment mechanism that fuels the emergence of AI agent hubs as a means of preserving competition in AI agent markets.

2. AI Agents

AI agents constitute algorithmic systems entrusted with autonomously formulating and implementing decisions on behalf of users. While autonomous

<https://www.cio.com/article/228353/big-data-is-not-a-barrier-to-entry.html#:~:text=Just%20as%20big%20data%20is,free%20data%20available%20to%20them.>

⁸ Michal S. Gal & Orla Lynskey, *Synthetic Data: Legal Implications of the Data-Generation Revolution*, 109 Iowa L. Rev. 1087 (2024).

⁹ Tucker & Wellford, *supra* note 7; Centre on Regulation in Europe (CERRE), *Big Data and Competition Policy - Market Power, Personalized Pricing and Advertising*, at 34 (February 16, 2017).

¹⁰ *Id.*, at 33; *Prediction Machines*, *supra* note 1, at 44.

¹¹ Stu Woo & Raffaele Huang, *How China's DeepSeek Outsmarted America*, Wall St. J. (Jan 28, 2025), <https://www.wsj.com/tech/ai/china-deepseek-ai-nvidia-openai-02bdbbce>; Gil Press, *DeepSeek Means The End Of Big Data, Not The End Of Nvidia*, Forbes (Jan 30, 2025), <https://www.forbes.com/sites/gilpress/2025/01/30/deepseek-means-the-end-of-big-data-not-the-end-of-nvidia/>. There are, however, allegations that DeepSeek unlawfully used data distillation to circumvent the full costs of collecting raw data. Ariel Zilber & Thomas Barrabi, *OpenAI Says It Has Proof DeepSeek Used Its Technology to Develop Its AI Model*, N.Y. Post (Jan 29, 2025), <https://nypost.com/2025/01/29/business/openai-says-it-has-proof-deepseek-used-its-technology-to-develop-ai-model/>.

algorithmic decision-making is not an entirely new phenomenon, the rise of AI agents substantially amplifies this capability, enabling agents to operate across a myriad of interconnected domains and to weigh a much larger set of parameters. AI agents can be programmed to execute complex, coordinated tasks—from optimizing manufacturing to handling procurement. The potential efficiency gains and improvements in decisional quality are so significant that all major players in AI markets are investing heavily in the development of such AI agents. This paradigm shift has even precipitated the emergence of a new term in the computer science literature to describe such algorithmic actions: agentism.

One potential function of such AI agents is their capacity to make purchasing decisions on behalf of consumers (what Gal and Elkin-Koren called Algorithmic Consumers)¹² or sale-related decisions on behalf of sellers. This transition toward algorithm-mediated commerce is already underway.¹³ The imminent commercial landscape will likely feature digital agents capable of managing complete transactional processes within e-commerce ecosystems. These agents predict consumer preferences, select appropriate products or services, conduct negotiations, and execute transactions—potentially eliminating human intervention from commercial decision processes. While buyers and sellers could theoretically develop these algorithms themselves, their complexity makes it more likely that specialized external providers will predominate in their supply.

3. Network Effects Driving Concentration in Markets for AI Agents

Network effects—typically associated with digital platforms—refer to the phenomenon in which a network’s value for its participants increases as more users join, as seen in the increased value of WhatsApp with each additional user.¹⁴ In the context of AI agents, delegating purchasing or selling authority to an AI system transfers associated bargaining power. Consequently, when numerous consumers authorize a single AI to handle their purchasing, they effectively endow it with collective negotiating leverage to withhold acquisitions or redirect their business elsewhere. Even if each consumer provides the AI with distinct parameters for decision-making, the AI agent inherently aggregates previously discrete purchasing decisions, thereby consolidating collective bargaining capacity.¹⁵ Such bargaining power consolidation can be the result of the aggregation of purchases of a similar product or service (such as flashlights). Crucially, this aggregation of bargaining power can also operate across multiple heterogeneous markets. For example, an AI

¹² Michal S. Gal & Nival Elkin-Koren, Algorithmic Consumers, 30 Harv. J.L. Tech. 309 (2017).

¹³ Id., 315-316; Ariel Ezrachi & Maurice E. Stucke, The Role of Secondary Algorithmic Tacit Collusion in Achieving Market Alignment, 26 Vand. J. Ent. & Tech 461, 465 (2024).

¹⁴ Raz Agranat, Uncovering the Role of Hubs: A Network Science Perspective on Platform Competition, 29 Stan. J. L. Bus. & Fin. 212 (2024).

¹⁵ Gal and Elkin-Koren, *supra* note 12.

agent managing purchases for consumers of both flashlights and tomatoes can leverage its cross-market bargaining power when negotiating with a retailer like Target, which sells both products. Consequently, AI agents with substantial user bases secure preferential terms, demonstrating superior performance metrics relative to competitors with limited user populations. This competitive advantage initiates a self-reinforcing positive feedback mechanism attracting additional users, potentially culminating in market foreclosure for competing AI agents with insufficient network scale.

The same dynamic applies to sellers. When numerous sellers entrust their selling decisions to a single AI agent, that AI gains substantial negotiating leverage over buyers. It can facilitate the alignment of commercial and pricing terms or withhold the aggregated offerings of distinct sellers from certain buyers. By securing better terms for its seller-users, the AI agent attracts more sellers and further strengthens its bargaining power, reinforcing its market position at the expense of rival AI agents. Again, this aggregation can extend across multiple markets: the AI agent can aggregate tomato and flashlight sales to create bargaining power in both markets.

Harrington and others have recognized that once a leading AI agent establishes its market position, a self-reinforcing feedback loop may progressively enhance its capacity to increase users' profitability,¹⁶ thereby intensifying the risk of market concentration and even tipping. Yet that risk has been primarily attributed to the AI agent's access to its users' confidential data rather than its ability to aggregate purchasing decisions vis-à-vis its users' trading partners.¹⁷ Building on the longstanding recognition that network effects can drive exclusionary outcomes—a concern traced back to the Microsoft cases of the 1990s¹⁸—we shift the focus to how dominant AI agents, by aggregating purchasing and selling decisions, can increasingly entrench their market position and foreclose rivals.

4. Lessons from Network Science

We can gain a deeper understanding of these dynamics through the emerging scientific discipline of network science, which studies the structure and properties of complex systems, including social, technological, biological, and economic ones. This discipline reveals that, despite their seemingly disparate nature, the topology and evolution of real-world networks are governed by similar organizing principles. Notably, we have employed network science in our previous work on digital platforms—both to analyze the robustness of the entry barriers arising from network

¹⁶ Joseph E. Harrington Jr., *The Challenges of Third Party Pricing Algorithms for Competition Law, Theoretical Inquiries L.* (forthcoming 2025).

¹⁷ Ezrachi & Stucke, *supra* note 13, at 479-481. For an exception see Gal & Elkin-Koren, *supra* note 12.

¹⁸ Raz Agranat & Michal S. Gal, *The Microsoft Formula for Network Effects: Groundbreaking Yet Stagnant, Antitrust L. J.* (forthcoming 2025).

effects and to illuminate the market power of influential entities operating on those platforms.¹⁹

Let us conceptualize the transition to an AI-agent-driven economy as an emerging network, where a growing number of trading partners engage through AI agents. Figures (1-3) illustrate the evolution of this network, focusing for simplicity on AI-driven commerce from the buyers’ perspective, with the understanding that similar dynamics can apply to sellers.

In the early phase, the use of AI-agents is limited to early adopters of the technology who form a minority of buyers.




Figure 1: Early stages of seller adoption in the AI Agent market

In network science terminology, the feedback loop described above indicates that the transition to AI-driven commerce is governed by two mechanisms: Growth and Preferential Attachment.²⁰ Growth means that the number of nodes in the network (in our drawing the green balls representing buyers) increases as more buyers (and sellers) adopt AI-agent technology. Preferential Attachment means that when new nodes are added to a network (that is, when more buyers start using AI agents for their purchases), they will tend to link to more connected nodes. In our case, buyers gravitate toward AI agents already used by other buyers, due to their greater bargaining power. Crucially, not only would new buyers connect to AI agents in accordance with preferential attachment, but existing buyers (i.e. who already use AI-agents) may switch AI agents in search of more favorable outcomes.²¹

¹⁹ Agranat, *supra* note 14 at 240-257; Raz Agranat & Michal S. Gal, Hub Power and Hub(uses): Power Dynamics in Platform Ecosystems(2025), <https://dx.doi.org/10.2139/ssrn.5136029>.

²⁰ Albert-László Barabási, *Network Science* 166-169 (2016).

²¹ This means that additional mechanisms beyond growth and preferential attachment shape the network’s topology; internal link addition and deletion, which involve changes to connections between existing nodes, also play a significant role. *Id.*, at 217-219.




Figure 2: Increasing seller adoption of AI agents

Over time, preferential attachment will likely propel the rise of hubs in the form of dominant AI agents. Rival AI agents that fail to attract sufficient sellers will gradually fall behind, and may ultimately be driven out of the market. As a result, market power would be concentrated in the hands of the largest AI agents, with their market position entrenched behind a network effects moat.




Figure 3: Preferential attachment elevates more-connected AI-agents at the expense of less-connected rivals

5. Regulatory Challenges

Antitrust law faces a significant challenge in addressing the aggregation of bargaining power by AI agents. First, the AI agent may lack the necessary market share to trigger monopolization/abuse of dominance prohibitions.²² Consider retailer Target dealing with an AI agent that represents 25% of the buyers. This percentage is well below the conventional threshold for triggering such prohibitions, yet Target may yield given that otherwise it will lose a large part of its sales. Furthermore, AI agents can aggregate the bargaining power of buyers operating across distinct markets. Consider an AI agent representing 20% of flashlight buyers, 15% of tomato buyers, and 10% of mug buyers. This cross-market aggregation will

²² Section 2 of the US Sherman Act, Section 102 of the EU treaty for the Functioning of the European Union (TFEU); See Ezrachi & Stucke, *supra* note 13, at 488.

allow the AI agent to exert pressure far beyond what these seemingly minor market shares suggest. It also makes assessing the AI agents’ market power incredibly complex. Furthermore, even if a monopolistic market position can be proven, the conduct itself might not be prohibited. Absent exclusionary conduct, the issue is one of exploitation, focusing on concerns of upward pricing pressure²³ or conversely, on disproportionate buyer power vis-à-vis sellers.²⁴ Yet even those jurisdictions that prohibit exploitation as such, do so only in extreme cases.

Alternatively, the contractual relationship between AI agents and their users potentially enables regulating the aggregation of bargaining power by AI agents through prohibitions on agreements in restraint of trade. The concern is not only that AI agent might enable coordination among their own users. Rather, as Ezrachi and Stucke argue, when only a handful of AI agents operate in a particular market, the agents might further engage in oligopolistic coordination among themselves.²⁵ Yet the application of prohibitions on coordinated conduct faces obstacles. Users of AI agents do not necessarily overtly collude but rather independently choose the most effective AI service in terms of securing favorable terms. Put differently, the AI agent facilitates tacit—and from the standpoint of its users, potentially inadvertent—coordination. One could consider applying the hub-and-spokes doctrine to the situation, with the AI agent as the hub and its users as spokes, but this approach has difficulties. Condemning the AI agents’ users when they are unaware of the aggregation, acting simply on the rational choice of selecting the most effective available service, seems inappropriate.²⁶ A hub-and-spokes framework might also fail to capture the aggregation of bargaining power across distinct markets, which may be substantial.

On a policy level, Harrington contends that condemning AI agents’ growth is problematic when their aggregation of users’ transactional decisions is a direct result of the substantial efficiencies they provide to their users.²⁷ Of course, if these efficiencies are not inherently dependent on such aggregation, AI agents could be

²³ Such pressure may stem from two sources: the coordination of sellers’ decision-making due to their use of a shared AI-agent, akin to a hub-and-spoke scheme; and either overt or tacit collusion between multiple AI-agents. Gal & Rubinfeld, *supra* note 4, at 694-696; Ezrachi & Stucke, *supra* note 13, at 468-473

²⁴ Gal & Elkin-Koren, *supra* note 12, at 332-334

²⁵ *Id.*

²⁶ The doctrinal challenge of condemning unwitting participants in a hub-and-spoke conspiracy is captured in the *Eturas* decision (Judgement of Jan. 21, 2016, *Eturas and Others*, C-74/14, EU:C:2016:42). In this case, travel agencies utilized the E-Turas system for booking management. Eturas, the system provider, sent messages intended to align the agencies’ discount offerings. However, some agencies claimed they should not be condemned, as they did not read these messages and were unaware of the attempted alignment. The court determined that Eturas’s mere dispatch of the messages was not sufficient to condemn the unaware agencies, necessitating additional evidence to prove their involvement.

²⁷ Harrington, *supra* note 16.

regulated without eliminating them. Yet isolating the part of the code which leads to aggregation from those parts that lead to efficiency may prove challenging. Absent an effective method to perform such a separation, regulators may face a stark choice: either permit aggregation or prohibit the use of AI agents altogether.

Furthermore, as Gal and Elkin-Koren argue, AI agents that aggregate buyers' purchasing decisions can counterbalance the market power of dominant sellers, thereby leading to better trade terms.²⁸ However, it is crucial to ensure that the benefits of such buyer power end up predominantly benefitting the buyers rather than being captured by AI agents, particularly as the latter gain market dominance. Indeed, antitrust law has long maintained that monopoly power should not be countered with monopsony power, and vice versa. The concentration of AI agents risks provoking competitive harms that far outweigh the benefits, including stifling innovation in AI technology, excluding sellers unable to withstand the buyer power wielded by dominant AI agents, driving up prices, deteriorating the quality of goods due to aggregated seller power, and consequently, creating overall deadweight loss.

6. Conclusion

Unfortunately, the window for effective antitrust intervention in the burgeoning AI agent market is narrowing. The more we allow preferential attachment to propel a few AI agents to behemoth status while excluding their rivals, reversing course will become exceedingly more difficult. Should we wish to thwart this outcome, one way forward involves disrupting the preferential attachment mechanism by using technological or legal means to forestall the consolidation of purchasing or selling power by any single AI agent in or across markets. Given the difficulty in enforcing these means, we should likewise endeavor to reduce artificial entry barriers facing potential entrants—such as incumbent-imposed limitations on access to critical inputs—thereby preventing network effects dynamics from being compounded by additional entry hurdles. At any rate, antitrust and consumer agencies would do well to shift their immediate attention to this looming threat of tipping in the AI agent market.

²⁸ Gal & Elkin-Koren, *supra* note 12.